PCT/AU00/00998 #2

AUSTRALIA

REC'D 13 SEP 2000
WIPO PCT

Patent Office
Canberra

4

I, KAY WARD, ACTING MANAGER EXAMINATION SUPPORT AND SALES hereby certify that annexed is a true copy of the Provisional specification in connection with Application No. PQ 2444 for a patent by NILS MARCHANT filed on 25 August 1999.

PRIORITY DOCUMENT
SUBMITTED OR TRANSMITTED IN COMPLIANCE WITH RULE 17.1(a) OR (b)

COMMONWEALTH OF AUSTRALIA
AUSTRALIA
PATENT OFFICE

WITNESS my hand this
First day of September 2000

K Ward

KAY WARD
ACTING MANAGER EXAMINATION
SUPPORT AND SALES

PRIORITY DOCUMENT
SUBMITTED OR TRANSMITTED IN COMPLIANCE WITH RULE 17.1(a) OR (b)

APPLICANT: **NILS MARCHANT**

NUMBER:

FILING DATE: **25 AUGUST 1999**

AUSTRALIA

PATENTS ACT 1990

PROVISIONAL SPECIFICATION

FOR THE INVENTION ENTITLED:

"A SWITCH AND A DISTRIBUTED NETWORK PROTECTION SWITCHING SYSTEM INCORPORATING SAID SWITCH"

The invention is described in the following statement:

A SWITCH AND A DISTRIBUTED NETWORK PROTECTION SWITCHING SYSTEM INCORPORATING SAID SWITCH

This invention relates to a switch and a distributed network protection switching system incorporating said switch, particularly, although not exclusively, for the protection of flows through networks against the failure of channels or sites in the network.

Networks are employed to convey flows between end users. The end users may be persons, computers, other machines, biological organs or other entities which are able to transmit or receive flows.. The events beyond the control of the network user cause networks to fail. Each failure within a network may disrupt one or more end users, in particular when the network is being relied upon for mission critical purposes. Network protection switching systems reduce the detrimental effects of failures upon end users. Some systems do so by switching flows away from failed parts of the network to healthier parts, if any exist. The failure and the protection switching action both lead to a period of disruption to the end user. A quick protection switching response time will reduce the disruption experienced by the end users.

Throughout this specification the term "network" is used in a generic sense to describe a set of two or more sites and one or more links that connect those sites together in any topology. A network supports the end to end transfer of flows between sites across a concatenation of one or more links within that network. Each link is unidirectional, has one source end, and has one or multiple destination ends. Each link transfers a flow or flows from the source end to one or more destination ends. A flow transmitted from a site onto a healthy link is transported to the destination site or sites. To form a bidirectional communication channel between two sites, links can be assembled as contra-flowing pairs.

It is important to note that the nature of the flow in one direction need not be the same as the flow in the opposite direction. Each site is able to transmit one or more

flows onto one or more links, and to receive flows from one or more links. Each link at each site is either an incoming link or an outgoing link depending on the direction of flow carried by that link. The receipt of any flow by a site from an incoming link may become unreliable while that link has failed. The transmission of a flow from a site may become unreliable when the site has failed.

Throughout this specification the word "flow" is intended to denote anything that can flow between two sites and includes the transfer of a communication, signal, substance, entity, energy, pressure, absence of pressure, vacuum, waves, particles or group of vehicles from one site to another site or sites. A single link can carry simultaneously one or more distinct and parallel flows. A single physical medium may carry distinct and opposing links or flows. Specific examples of flows include electrical, radio or optical signals, whether in digital or analogue form that convey information between sites within the network; hydraulic pressure conveyed along tubes, pipes or hoses; gas, liquid, slurry, solids, people, other beings or organic matter; pneumatically propelled canisters; a vacuum propagated along a tube, pipe or hose; or a transfer of energy.

One of the difficulties against which network engineers and protocol architects have been struggling is that of insuring that information about a network failure is conveyed to the required sites in order to initiate protection switching. The problem is to carry such information through a faulty network. It is not always straight forward to determine which part or parts of a network can be relied upon to carry this information. The information about the failures should be carried on channels or links which have not failed. Sometimes even this cannot be achieved. There have been many proposals to protect networks against failures but these all have one or more of the following disadvantages:

- They require a separate network in addition to the protected network on which to transmit fault monitoring and control signals.
- The devices which make and implement protection switching decisions are installed at sites which are geographically or topologically remote from the links

or sites which they protect. In order to switch flows away from failed parts of the network, they rely upon the transmission through that network of signals which carry information about the failure. Systems which rely upon the transmission of such fault information signals have the following disadvantages:

- It is not always possible to find a path which has not failed for the fault signals. It is thus not always possible to guarantee that the fault information signals are able to reach their destinations.
- Network users may suffer unnecessarily long disruptions because switching decisions and operations are not performed near to the points of failure and because of the time required for: i) fault signals to be generated; ii) information about a fault to be conveyed across part of a network; iii) that information to be processed; iv) protection switching decisions to be made; v) control signals to be conveyed across part of a network in order to implement switching decisions; and vi) the switching decisions to be executed.
- The network may become flooded with a large number of fault messages which may overload the network and reduce the availability of the network to the end users.

- They rely upon central sites which themselves are vulnerable to failure.
- The time required to detect that a fault has occurred may be long. The decision that a link or site has failed is delayed until after a flow has become absent for at least a specified "time-out" period, resulting in a delay before protection switching may be initiated.
- The time required for protection switching to take effect is proportional to the number of sites in a network.
- They do not protect networks whose topology is more complex than a simple point to point link.
- They require inflexible and costly non-optimal network topologies to be installed.
- They require looped or ring topology networks.
- They do not protect against the failure of sites.
- They do not permit more than one redundant link to protect each protected link.
- They do not support an arbitrary level of redundancy.

- They require redundant links to be installed where they are not required.
- In communications networks they do not operate independently of and transparently to open systems interconnection (OSI) layer services or protocols above the physical medium dependent (PMD) sub-layer of the physical layer (layer one). Separate protection switching devices must be constructed for each type of protected network.
- They degrade network performance while a fault condition exists
- They reduce network capacity while a fault condition exists.
- They require the unnecessary duplication or replication and simultaneous transmission of flows at times when no fault exists.

It is an object of the present invention to provide a switch and a distributed network protection switching system that enables end to end connectivity across a network to be maintained in a highly reliable fashion by switching flows away from failed parts of the network.

According to a first aspect of this invention there is provided a switch comprising at least:

first, second and third ports (P_1-P_3) each having an input and an output;

the switch arranged so that a first flow presented to the input of one of P_1-P_3 is delivered to the output of an other of P_1-P_3, and a second flow presented to the input of said other of P_1-P_3 is delivered to the output of said one of said P_1-P_3;

detection means for detecting a predetermined characteristic of the flows presented at the input of each of P_1-P_3; and,

control means which, upon the detecting means detecting said predetermined characteristics in one of said first flow and said second flow,internally diverts the other of the first flow and second flow to be presented to the output of the remaining one of P_1-P_3.

Preferably said switch further includes timer means for counting a time T for which the detecting means detects the existence of said predetermined characteristic of the flows and wherein said control means only diverts the other of the first and second

flows to the output of said remaining one of P_1-P_3 when the time $T \leq$ a predetermined time T_{wait}.

Preferably said switch includes a dummy flow means for producing a dummy flow for the period when $T < T_{wait}$ and, during this period, said control means delivers said dummy flow to the output of the port to which said one of said first and second flows would be delivered in the absence of that flow being detected as having said predetermined characteristic.

In one embodiment, said predetermined characteristic is the absence of said flow for said period T_{wait}. In one embodiment, the period T_{wait} is 0.

In an alternate embodiment, said predetermined characteristic is a predetermined reduction in the rate of flow at said inputs. In an alternate embodiment when said flows relate to communication signals, said predetermined characteristic is a predetermined bit error rate, or signal to noise ratio.

In one embodiment, said dummy flow means is in the form of a generator for generating a flow of the same type as the flow presented to the inputs of said switch.

In an alternate embodiment, said dummy flow means includes means for sampling and subsequently replicating the flow presented to the inputs of the switch.

Preferably said detection means is further able to detect the absence of said predetermined characteristic after said control means has internally diverted said other of the first flow and second flow to the output of the remaining one of P_1-P_3, whereup said control means rediverts said other of the first flow and second flow to be presented to the output of the other or one of P_1-P_3, as the case may be.

According to a further aspect of the present invention there is provided a distributed network protection switching system for a network having at least first and second sites (X_1, X_2), a first channel to allow bidirectional transfer of flows between said

sites and, at least one further channel to provide an alternate route for bidirectional transfer of flows between said sites, each channel having a unidirectional incoming link and an unidirectional outgoing link, the system comprising at least:

a first switch (S_1) and a second switch (S_2), each of S_1, and S_2 being in accordance with the first aspect of this invention;

S1 coupled to the first site (X_1) so that a flow out of X_1 is presented to the input of any one of P_1-P_3 of S_1 and a flow into X_1 is delivered from the output of said one of P_1-P_3 of S_1;

S_2 coupled to X_2 so that a flow out of X_2 is presented to the input of any one of P_1-P_3 of S_2 and a flow into X_2 is delivered from the output of said one of P_1-P_3 of S_2;

the outgoing link of channel C_1, viewed from X_1, connected between the output of an other of P_1-P_3 of switch S_1 and the input of an other of P_1-P_3 of switch S_2;

the incoming link of channel C_1, viewed from site X_1, connected between the input of said other of P_1-P_3 of switch S_1 and the output of said other of P_1-P_3 of switch S_2;

the outgoing link of channel C_2, viewed from site X_1, connected between the output of the remaining one of P_1-P_3 of switch S_1 and the input of the remaining one of P_1-P_3 of switch S_2 and, the incoming link of channel C_2 viewed from site X_1 being connected between the input of the remaining one of ports P_1-P_3 of switch S_1 and the output of the remaining one of P_1-P_3 of switch S_2;

whereby, in use, upon said detection means of one of S_1 and S_2 detecting a predetermined characteristic of the input for channel C_1 internally diverts the flow directed to the output of the port containing that input to the output of the remaining port thereby causing the detection means of the other one of S_1 and S_2 to detect the absence of the flow at the input of the other one of switches S_1 and S_2 for channel C_1 so that the flow delivered to the output of the other port of switch S_2 is directed to the output of the remaining port of switch S_2 thereby switching the channel of communication between the first and second sites X_1,X_2 from channel C_1 to channel C_2.

Embodiments of the present invention will now be described in detail with reference to the accompanying drawings in which:

Figure 1 is a block diagram of a switch in accordance with an embodiment of the present invention;

Figures 2A, 2B & 2C depict three internal switching configurations of the switch shown in Figure 1;

Figure 3 is a block diagram of the switch depicted in Figure 1;

Figures 4-14 depict the switching configurations employed in 11 different states of one state machine used to describe the operation of the switch shown in Figure 1;

Figure 15 illustrates an embodiment of one network incorporating the switch and distributed network protection switching system in accordance with the present invention;

Figure 16 illustrates a distributed network protection switching system for protecting a network against a site failure;

Figure 17 illustrates a distributed network protection switching system for protecting a network against a point to point channel failure;

Figure 18 illustrates a distributed network protection switching system for protecting a network against a link failure between switches in the network;

Figure 19 illustrates on possible set of timing relationships in the operation of the switch; and,

Figure 20 depicts a bypass device that can be incorporated in or associated with the switch.

Referring to the accompanying drawings, switch 10 in accordance with an embodiment of this invention is provided with three ports P_1-P_3. Each port has an input I and an output O. The input of port 1(IP_1) and the output of port 1 (OP_1) are designated by reference numbers 12 and 14 respectively. Similarly IP_2, OP_2, IP_3, OP_3, are designated as reference numerals 16, 18, 20 and 22 respectively. The switch 10 has three incoming links 24i, 26i and 28i for presenting flows to the inputs 12, 16, and 20 respectively and three outgoing links 24o, 26o, and 28o for carrying flows delivered from the outputs 14, 18 and 22 respectively. The pairs of links at each port can be considered to form a bidirectional channels 24, 26 and 28.

The switch 10 is arranged so that a first flow presented to the input of any one of ports P_1-P_3 (for example a flow on link 24i to input 12) is delivered to the output of an other of P_1-P_3 (ie output 18 of port P_2) and a second flow presented to the input I of that other of P_1-P_3 (ie in this case the flow on link 26i to the input 16 of port P_2) is delivered to the output O of said one of ports P_1-P_3 (ie output 14 of port P_1). Thus, the channel 24 is connected to the channel 26 via internal paths in switch 10 connecting input 12 to output 18, and output 14 to input 16.

Detection means (not shown in Figure 1) is included in the switch 10 for detecting a predetermined characteristic of the flows presented to the inputs of each of the ports P_1-P_3. Associated with the detection means is a control means (not shown in Figure 1) that, upon the detecting means detecting the predetermined characteristic in one of the flows presented to the inputs 12 or 16 internally diverts the flow from the other one of these inputs to the output 22 of the remaining port P_3.

Depending on the state of the switch 10, flows through the switch 10 can be passed through different combinations of pairs of ports P_1-P_3. Figure 2A illustrates the switch 10 in a configuration where flows pass through ports P_1 and P_2 of the switch 10. In Figure 2B, the flows are directed through ports P_1 and P_3; and in Figure 2C

the flows are directed through ports P_2 and P_3. The configuration taken by the switch 10 is dependent upon the characteristics detected at each of the inputs of all three ports as well as the current state of the switch 10.

Referring to Figure 3, the switch 10 is provided with flow detectors 30, 32 and 34 connected with the inputs 12, 16 and 20 respectively for detecting a predetermined characteristic of the flows presented to those inputs by the incoming links 24i,26i and 28i. As explained in greater detail hereinafter, the predetermined characteristic may be the simple presence or absence of a flow, or a predetermined bit error rate (for a digital transmission) or a signal to noise ratio (for an analogue transmission). However for the purpose of the following discussion, the detectors 30, 32 and 34 are structured to detect the present or absence of the flow, including a dummy flow, at each of the incoming links 24i, 26i and 28i. The presence of a flow is interpreted as "link healthy". The absence of a flow for greater than a time T_{detect} (T_{detect} is described below) is interpreted as "link not healthy" or "link failed". Each of the detectors 30, 32 and 34 is coupled by a corresponding line to a controller 36. The flow on incoming link 24i is passed to internal switch 38 for directing the flow along one of alternate paths 40,42. The flow on incoming path 26i can be switched via internal switch 44 along alternate paths 46,48. The flow on incoming line 28i is provided to switch 50 for switching between alternate paths 52 and 54. The switch 10 includes internal output switch 56 leading to output 14; output switch 58 leading to output 18; and output switch 60 leading to output 22. Switch 10 further includes a dummy signal means 62 for generating a dummy flow or sampling and replicating a flow on one of the incoming links 24i, 26i, 28i. The output of the dummy generator is presented to switch 64. The switch 64 is linked to switch 60 by path 66, linked to switch 56 by path 68, and linked to switch 58 by path 70. Each of the switches 38, 44, 50, 56, 58, 60 and 64 has an associated control block 72. Each control block 72 receives control signals from the controller 36 via respective signal lines 74.

Also included is a power supply block 76 for providing a regulated power supply to the internal components of switch 10.

The purpose of the dummy flow produced by the dummy generator 62 is to enable one or more switches 10 elsewhere in the network to detect the presence of that dummy flow and to interpret the detection that flow as "incoming link healthy". A dummy flow from one switch 10 is able to stimulate switching decisions at another switch 10. The requirements of the dummy flow and the dummy flow source which generated are not stringent. For example, in an optical embodiment of the switch 10, a continuous stream of light of suitable wave length would suffice as a dummy flow. Where optical receiver power specifications or safety requirements are to be met (for example to avoid eye damage) a low duty cycle rectangular wave form would be employed to reduce the transmitted optical power. In an electrical embodiment, a simple square wave or rectangular wave would suffice. Where a periodic wave form is used the longest off period should be less than a predetermined time T_{detect}.

Optionally, instead of generating a dummy flow, in another embodiment of the invention the dummy flow means would replicate a selected incoming flow so that a copy of that incoming flow could be transmitted onto a selected set of outgoing lengths.

The controller 36 monitors the presence or absence of flows detected at each of the incoming links 24i, 26i and 28i via corresponding flow detectors 30, 32 and 34. In this configuration it can be seen that the flow on incoming link 24i can be delivered to the output 18 via switch 38, path 40 and switch 58; or output 22 via switch 38, path 42, and switch 60. The flow and input 16 carried on incoming line 26i can be delivered to the output 14 via switch 44, path 46 and switch 56; or output 22 via switch 44, path 48 and switch 60. A flow incoming line 28i at input 20 can be delivered to the output 14 via switch 50, path 52 and switch 56; or to the output 18 via switch 50, path 54 and switch 58. The dummy flow provided by dummy generator 62 can be delivered to output 14 via switch 64, path 68 and switch 56; to output 18 via switch 64, path 70 and switch 58 or to output 22 via switch 64, path 66 and switch 60.

By supplying the appropriate combination of control signals to these internal switches, the controller 36 is able to cause a flow presented at the input of one port to be directed to the output to either one of the remaining ports. The controller 36 is also able to cause the dummy flow to be switched to any one of the three output ports, 14, 18 and 22. Additionally, the controller 36 is able to sense whether adequate power supply is available from the power supply box 76. In the absence of adequate power, or if no flow is detected on both normal incoming links 24i and 26i, the control block 36 activates two nominal normal paths to allow a flow presented at input 12 to flow to output 18 and a flow presented at input 16 to flow to output 14.

It is possible to implement the controller 36, dummy flow generator 62 and switches 38, 44, 50, 56, 58, 60 and 64 in a field programmable gate array (FPGA) such as the **ALTERA** EPF8282V which performs a switching algorithm based on a state machine described hereinafter. The input and output at each port can be realised by the use of a transceiver such as the Crystal Technologies CS61304A T1E1 line interface. This would provide signal bit rates at 1.544Mbps and 2.048Mbps. This component extracts bit clock and line data from an incoming flow, and generates a "loss of signal" (LOS) output signal when the absence of an incoming flow is detected. The data, clock and LOS outputs are supplied to the FPGA. A local timing source can be provided by a crystal such as the 32.768Mhz FOX F5C HS oscillator. The FPGA generates a dummy flow derived from the local clock and in accordance with the state machine (to be described below). Incoming and dummy data and clock signals are supplied from the FPGA to the line interfaces which transmit them to the outgoing links. Local and remote management interfaces enable the switch 10 to provide status information and to be controlled.

The operation of the switch 10 may be specified by a state machine. Figures 4-14 illustrate the 11 possible states in which the switch 10 would be found.

The 11 states are: OFF (O); NORMAL (N); IDLE (I); WAIT A (WA); WAIT B (WB); SKEW A (SA); SKEW B (SB); LOCK A (LA); LOCK B (LB); HUNT A (HA); and HUNT B (HB). Each is described below.

OFF (O) (Figure 4)

No power is supplied to the switch 10. If a by-pass device is installed it operates to by-pass the switch 10. Each of inputs 12 and 16 is through enabled if possible to outputs 18 and 14 respectively. The broken lines represent that. Any flows which might be present at links 24i and 26i are thus transmitted to links 26o and 24o respectively. The switch 10 remains in the OFF state until power is supplied to it. It does not perform any active functions. Should power be applied the switch 10 switches to one of the other ten states whose descriptions follow below. The switching decision is not affected by the presence or the absence or nature of any flow at 28i, 24i or 26i. No flow is transmitted from 28o. No dummy flow is transmitted.

NORMAL (N) (Figure 5)

A flow is detected at inputs 12 and 16 for incoming links 24i and 26i. The switch 10 is through enabled. Each flow from incoming links 24i and 26i is transmitted onto the normal outgoing links 26o and 24o respectively via outputs 18 and 14. The switching decision is not affected by the presence or absence or nature of any flow on link 28i. No flow is transmitted from output 22 onto link 28o. No dummy flow is transmitted.

IDLE (I) (Figure 6)

No flow is detected at either inputs 12 or 16 for links 24i or 26i. The switch 10 is through enabled. No flows are transmitted from outputs 14 or 18 onto links 26o or 24o. The switching decision is not affected by the presence or absence or nature of any flow presented to input 20 on link 28i. No flow is transmitted from output 22 onto link 28o. No dummy flow is transmitted.

WAIT A (WA) (Figure 7)

No flow is detected at input 16 for link 26i. A flow is detected at input 12 on link 24i. Link 24i is through enabled so that any flow at link 24i is transmitted from link 26o. The dummy flow 62 is active and transmitted from 24o. A timer TB is timing the delay period T_{WAIT}. (T_{WAIT} is described below.) The switching decision is not

affected by the presence or absence or nature of any flow at link 28i. No flow is transmitted from link 28o.

WAIT B (WB) (Figure 8)

No flow is detected at input 12 for link 24i. A flow is detected at input 16 on link 26i. The flow on link 26i is through enabled to be delivered to output 14 and thus be transmitted on link 24o. The dummy flow 62 is active and delivered to output 18 to flow on link 26o. A timer TA is timing the delay period T_{WAIT}. (T_{WAIT} is described below.) The switch 10 is not affected by the presence or absence or nature of any flow on link 28i. No flow is transmitted from 28o.

SKEW A (SA) (Figure 9)

A flow on link 24i is detected at input 12. No flow is detected at input 16 for link 26i. The link 24i flow is delivered to output 22 to be transmitted on the alternative outgoing link 28i. The dummy flow 62 is active and delivered to output 14 to be transmitted on link 24o. No flow is transmitted on link 26o.

SKEW B (SB) (Figure 10)

A flow on link 26i is detected at input 18. No flow is detected at input 12 for link 24i. The link 26i flow is delivered to output 22 to be transmitted on alternative outgoing link 28o. The dummy flow 62 is active and transmitted from link 26o. No flow is transmitted from link 24o.

LOCK A (LA) (Figure 11)

A flow is detected at each of inputs 12 and 20 for incoming links 24i and 28i. No flow is detected at input 16 for link 26i. The incoming flow from link 24i is transmitted onto link 28o. The incoming flow from link 28i is transmitted onto link 24o. No flow is transmitted from link 26o. No dummy flow is transmitted.

LOCK B (LB) (Figure 12)

A flow on each of incoming links 26i and 28i is detected at inputs 18 and 20 respectively. No flow is detected at input 12 for link 24i. The incoming flow from

link 26i is transmitted onto link 28o. The incoming flow from link 28i is transmitted onto link 26o. No flow is transmitted from link 24o . No dummy flow is transmitted.

<u>HUNT A (HA) (Figure 13)</u>

No flow is detected at input 18 for link 26i. A flow is detected on link 24i. The flow received on link 24i is through enabled and transmitted from link 26o. The dummy flow 62 is active and transmitted from link 28o. No flow is detected at link 28i. No flow is transmitted from link 24o.

<u>HUNT B (HB) (Figure 14)</u>

No flow is detected at input 12 for link 24i. A flow is detected on link 26i. The flow received on link 26i is through enabled and transmitted from 24o. The dummy flow 62 is active and transmitted from link 28o. No flow is detected at link 28i. No flow is transmitted from link 26o.

Transitions between the above described states are caused by the following events:

E1	**Power On**	The switch 10 commences to receive an adequate power supply.
E2	**A Found**	A flow is detected at link 24i.
E3	**A Lost**	A flow at link 24i ceases to be detected.
E4	**B Found**	A flow is detected at link 26i.
E5	**B Lost**	A flow at link 26i ceases to be detected.
E6	**Alt Found**	A flow is detected at link 28i.
E7	**Alt Lost**	A flow at link 28i ceases to be detected.
E8	**TA Timeout/ALT**	The T_{WAIT} timer TA expires while a flow is detected at link 28i.
E9	**TA Timeout/NALT**	The T_{WAIT} timer TA expires while no flow is detected at link 28i.
E10	**TB Timeout/ALT**	The T_{WAIT} timer TB expires while a flow is detected at link 28i.

E11 **TB Timeout/NALT** The T_{WAIT} timer TB expires while no flow is detected at link 28i.

E12 **Power Off** The switch 10 ceases to receive an adequate power supply. The switch 10 returns to the OFF state after this event, irrespective of any prior state.

The algorithm implemented at each switch 10, and the set of states, events and transitions between states which are employed here to describe the invention are not unique. The invention may be performed with alternative state machines and algorithms. Some state machines may employ a different number of state transitions than described here. For example the invention may also be performed using the state machine described here and with the following additional state transitions: IDLE to NORMAL; WAIT A to LOCK A; and WAIT B to LOCK B.

The switch 10 contains two delay timer functions TA and TB. TA is activated when the event "A Lost" occurs; TB is activated when the event "B Lost" occurs. Both TA and TB are configured to measure the delay time T_{WAIT}. Expiry of the timer TA after delay T_{WAIT} causes the switch 10 to switch from the WAIT B state to either the SKEW B or LOCK B state, depending on whether a flow is detected at input 20 for link 28i. Expiry of the timer TB causes the switch 10 to switch from the WAIT A state to either the SKEW A or LOCK A state, depending on whether a flow is detected at link 28i. While the switch 10 is in the WAIT A state, if the flow at link 24i is also subsequently lost the switch 10 remains in the WAIT A state until either: the timer TB expires; or the flow at link 26i is newly detected.

Optionally, the delay time T_{WAIT} may be set to zero in some switches 10. In such cases the transition through the WAIT states will be instantaneous.

One switch 10 state transition table is shown in table one below. While the switch is in any of the states shown at the left of the table, the occurrence of any of the events shown at the top of the table causes the switch 10 to change to the new state shown within the table. The Power On event only occurs while in the OFF state and is the

only event possible in the OFF state. Blank entries indicate inapplicable state transitions.

Present State ↓	Events											
	E1 Power On	E2 A Found	E3 A Lost	E4 B Found	E5 B Lost	E6 Alt Found	E7 Alt Lost	E8 TA Timeout / ALT	E9 TA Timeout / NALT	E10 TB Timeout / ALT	E11 TB Timeout / NALT	E12 Power Off
OFF	I											
I		HA		HB		I	I					OFF
HA			I	N		LA						OFF
HB		N			I	LB						OFF
N			WB		WA	N	N					OFF
WA		WA	WA	N		WA	WA			LA	SA	OFF
WB		N		WB	WB	WB	WB	LB	SB			OFF
SA			I	N		LA						OFF
SB		N			I	LB						OFF
LA			I	N			SA					OFF
LB		N			I		SB					OFF

Table 1: Network Repair Module State Transitions

The interaction between switches 10 and the operation of the state machine at each switch 10 is now described in detail by way of example. Two cases are examined: link failure; and site failure.

Figure 15 shows nine switches 10_A-10_I that are arranged to protect a bidirectional channel C_1 comprising a pair of links marked 78 in one direction and 80 in the opposite direction. The links 78 and 80 through switches 10_A-10_D form the normal channel which would be used when there are no faults. There are two redundant paths or channels: the first redundant channel C_2 comprises links 82,84 through

switches 10_E-10_G; and the second redundant channel C_3 comprises links 86 and 88 through switches 10_H and 10_I. The switches 10_A and 10_B would be at one site and the Switches 10_C and 10_D would be at a separate site.

The sequence of states in which each switch 10 finds itself during protection switching after a link failure is shown in table two below. The description and the table illustrate only one of many possible sequences of state transitions which may occur as a result of the failure. The description and the table do not illustrate all state changes which may occur.

Switches 10_A-10_D each commence in the NORMAL state (N) and carry flows along links 78,80. No flow is carried on the redundant channels C_2,C_3 and Switches 10_E-10_I commence in the IDLE state (I).

Let a failure occur in channel C_I between switches 10_B and 10_C. It does not matter whether the failure occurs only on a single link 78 or 80, or on both links 78 and 80, or if there is a time delay between a failure on each; the present invention will operate in all cases. In this example we shall assume that the failure occurs either on link 80 only, or on both links 78 and 80. The failure causes switch 10_B to detect the absence of a flow at its input 16, where the flow from switch 10_C via link 80 normally would be received. The absence of that flow causes switch 10_B to enter the WAIT A state (WA) and to present the dummy flow to its output 14 onto link 80 towards switch 10_A. The dummy flow is detected by switch 10_A and further downstream on link 80. If for some reason the link is repaired and the incoming flow on link 80 at input 16 of switch 10_B is again detected then switch 10_B returns to the NORMAL state. However, if the flow remains absent for at least a time T_{WAIT} (which would cause the timer TB to expire and the TB Timeout event to occur) then switch 10_B enters the SKEW A state (SA), where it ceases to deliver any flow to output 18 and thus link 78 ceases to transmit any flow towards switch 10_C. Instead, the incoming flow at input 12 of switch 10_B is switched onto the alternative outgoing link 82 via its output 22 and travels towards switch 10_E.

Switch	10_A	10_B	10_C	10_D	10_E	10_F	10_G	10_H	10_I
Initial State	N	N	N	N	I	I	I	I	I
After failure	N	WA	N	N	I	I	I	I	I
After T_{WAIT} at 10_B	N	SA	N	N	I	I	I	I	I
Flow link 82 reaches 10_E	N	SA	N	N	HA	I	I	I	I
Flow link 82 reaches 10_F	N	SA	N	N	HA	HA	I	I	I
Flow link 82 reaches 10_G	N	SA	N	N	HA	HA	HA	I	I
Loss detected at 10_C	N	SA	WB	N	HA	HA	HA	I	I
After T_{WAIT} at 10_C	N	SA	SB	N	HA	HA	HA	I	I
Flow link 84 reaches 10_G	N	SA	SB	N	HA	HA	N	I	I
Flow link 84 reaches 10_F	N	SA	SB	N	HA	N	N	I	I
Flow link 84 reaches 10_E	N	SA	SB	N	N	N	N	I	I
Flow link 84 reaches 10_B	N	LA	SB	N	N	N	N	I	I
Flow link 82 reaches 10_C	N	LA	LB	N	N	N	N	I	I
Final state	N	LA	LB	N	N	N	N	I	I

Table 2: Summary of Significant State Transitions in Response to Link Failure.

Note 1: The sequence of transitions shown above is only one of many valid possible sequences.
Note 2: No time scale is implied by Table 2.

Switch 10_E in the IDLE state, detects at its input 12 via link 82 the flow from output 22 of switch 10_B. That event causes switch 10_E to enter the HUNT A state (HA) in which it: i) delivers a dummy flow to its P_3 output 22; and ii) transmits the incoming flow it receives at its input 12 onto link 82 via its P_2 output 18. Switches 10_F and 10_G subsequently change from the IDLE state to the HUNT A state as a result of detecting the flow propagated to them along link 82 from switch 10_B.

After the disappearance of the flow on link 78 and thus at input 12 of P_1 of switch 10_C, switch 10_C enters the WAIT B state (WB) and then after a time delay T_{WAIT} enters the SKEW B state (SB) if it has not yet detected any flow at input 20 of P_3 via link 82. A dummy flow is transmitted by switch 10_C downstream along link 78 to switch 10_D and beyond. In the SKEW B state the incoming flow on link 80 at input 16 of P_2 of switch 10_C is switched to output 22 of $P_3(10_C)$. After travelling along link 84 that flow is detected at input 16 P_2 (10_G). Switch 10_G then either: if no flow is detected at input 12, P_1 (10_G), changes from the IDLE state to the HUNT B state (HB); or, (as shown in Table 2) if it has already detected a flow at

input 12, P_1 (10_G), changes from the HUNT A state to the NORMAL state. Eventually switch 10_G detects a flow at both input 12,P_1 and input 16,P_2 and enters the NORMAL state. The flow from output 14, P_1 (10_G) propagates further downstream along link 84 through switches 10_F and 10_E in that order, eventually causing these switches to switch into the NORMAL state.

The flow from switch 10_C via link 84 and switches 10_G, 10_F, 10_E is detected at input 20, P_3 of switch 10_B. That event causes switch 10_B to change from the SKEW A state to the LOCK A state (LA). The flow propagated along the link 82 causes switch 10_C to change from the SKEW B state to the LOCK B state (LB) when it detects the flow at input 20,P_3 switch 10_G. A complete bidirectional channel C_2 via switches 10_A, 10_B, 10_E, 10_F, 10_G, 10_C and 10_D now operates around the failed normal channel C_1. A subsequent failure of the redundant links 82 or 84 in channel C_2 causes a similar sequence of events and state changes at switches 10_A, 10_H, 10_I and 10_D which result in diversion of flows along channel C_3 via switches 10_A, 10_H, 10_I and 10_D.

Figure 16 shows nine switches 10_A-10_I that are arranged to protect a network against the failure of a site 90. Site 90 contains a set of one or more nodes 92 and two switches 10_B, 10_C. Table 3 shows the sequence of states after the failure of site 90. Switches 10_A-10_D commence in the NORMAL state (N), and switches 10_E-10_I commence in the IDLE state (I). If site 90 ceases to transmit any flows for any reason then no flows are transmitted on any of the four outgoing links 11, 13, 15 and 17 which emanate from site 90. Switch 10_A detects the loss of flow at its input 16 and switch 10_D detects the loss of flow at its input 12. The switch 10_A changes from the NORMAL state to the WAIT A state (WA) and switch 10_D changes from the NORMAL state to the WAIT B state (WB). Dummy flows are transmitted separately by switch 10_A onto link 19 and by switch 10_D onto link 21.

After a time delay T_{WAIT} at switch 10_A, this switch enters the SKEW A (SA) state. After a time delay T_{WAIT} at switch 10_D, this switch enters the SKEW B (SB) state. The flow on incoming link 23 to switch 10_A is transmitted onto link 25 coupled to

output 22 of switch 10_A; and the flow on incoming link 27 to switch 10_D is transmitted onto link 29 coupled to output 22 of switch 10_D. Switches 10_E, 10_G, 10_H switch in that sequence from the IDLE state to the HUNT A state (HA) as the flow is propagated along the link 25 from switch 10_A. Similarly, as the flow from output 22 of switch 10_D is propagated along link 29, switches 10_H, 10_I change from the IDLE state to the HUNT B state (HB) in that sequence. It is a combination of both the topology shown and the switching algorithms that ensures that switches 10_G and 10_H make contact and receive flows from each other.

When switch 10_G detects the flow at its input 20 from link 31, it changes from the HUNT A state to the LOCK A state (LA). Switch 10_G detects flows at input 12 from link 25 and at input 20 from link 31. When switch 10_H detects the flow on link 33 from switch 10G at input 20 it changes from the HUNT B state to the LOCK B state (LB). The flow received by switch 10_G at its input 20 is switched onto its output 14 and propagated on link 35 via switches 10_F and 10_E to switch 10_A. That flow causes switches 10_F and 10_G each to change from the HUNT A state to the NORMAL state in that sequence. When the flow is detected at input 20 of switch 10_A, that switch changes from the SKEW A state to the LOCK A state. Similarly, the flow received by switch 10_H at its input 20 is switched onto its output 18 and propagated by link 17 to switch 10_D via switch 10_I. That causes switch 10_I to change from the HUNT B state to the NORMAL state. When the flow is detected at input 20 of switch 10_D, that switch changes from the SKEW B state to the LOCK B state.

The switches have now constructed a bidirectional channel around site 90 via switches 10_A, 10_E, 10_F, 10_G, 10_H, 10_I and 10_D. The flow on incoming link 23 to input 12 of switch 10_A is carried via links 25, 33 and 17 to switch 10_D and is transmitted by switch 10_D onto link 21. The flow on incoming link 27 presented to input 16 of switch 10_D is carried via links 29, 31 and 35 to switch 10_A and is delivered by switch 10_A onto its output 14 to flow on outgoing link 19. The failed site 90 is thus by-passed.

Switch	10_A	10_B	10_C	10_D	10_E	10_F	10_G	10_H	10_I
Initial State	N	N	N	N	I	I	I	I	I
After site failure	WA	N	N	WB	I	I	I	I	I
After T_{WAIT} delays	SA	N	N	SB	I	I	I	I	I
Flow reach 10_E and 10_I	SA	N	N	SB	HA	I	I	I	HB
Flow reach 10_F and 10_H	SA	N	N	SB	HA	HA	I	HB	HB
Flow reaches 10_G	SA	N	N	SB	HA	HA	HA	HB	HB
Flow from 10_H reaches 10_G	SA	N	N	SB	HA	HA	LA	HB	HB
Flow from 10_G reaches 10_H	SA	N	N	SB	HA	HA	LA	LB	HB
Flow from 10_H reaches 10_F	SA	N	N	SB	HA	N	LA	LB	HB
Flow from 10_H reaches 10_E	SA	N	N	SB	N	N	LA	LB	HB
Flow from 10_H reaches 10_A	LA	N	N	SB	N	N	LA	LB	HB
Flow from 10_G reaches 10_I	LA	N	N	SB	N	N	LA	LB	N
Flow from 10_G reaches 10_D	LA	N	N	LB	N	N	LA	LB	N
Final state	LA	N	N	LB	N	N	LA	LB	N

Table 3: Summary of Significant State Transitions in Response to Site Failure.

Note 1: The sequence of transitions shown above is only one of many valid possible sequences.
Note 2: No time scale is implied by Table 3.

Site 90 is arranged in such a manner that a power failure at any switch 10_A-10_I or the node causes the power supply to all switches at site 90 to be switched OFF.

Note that the topology shown above is suitable to protect against the failure of a site, the failure of one or more links to the site, or combinations of failures of links and the site.

Switches 10 are able to protect simple point to point link topologies. The time delay T_{WAIT} may be set optionally to zero at the innermost two switches which are immediately adjacent to a protected link. Two benefits are provided by setting the time delay T_{WAIT} to zero:

- very high speed protection switching performance of the switches 10 may be achieved; and
- the design and manufacture of switch 10 may be simplified.

Figure 17 shows one topology which would be used to protect a point to point channel C_1 comprising pair of links 37 between a site 94 and a site 96. Two redundant channels C_2 and C_3 are shown. When channel C_1 represents a subscriber access connection to a public telecommunications network the site 94 would represent the subscriber's premises and the phantom line 39 represents the subscriber interface which would be crossed by the subscriber access links and redundant links. The equipment 98 would represent customer premise equipment (CPE) or other end user equipment. The site 96 would represent the public network. (Alternately when channel C_1 represents a trunk the sites 94 and 96 would represent separate exchange buildings or central offices.) Redundant channel C_2 is connected to the protected links at switches 10_B and 10_D. Redundant channel C_3 connected to the protected links at switches 10_A and 10_E. Further switches 10_C and 10_F are shown which provide additional internal protection to the network.

The innermost switches 10_B and 10_D which connect to the first redundant channel C_2 are the first switches to operate after a failure of any link 37 of channel C_1. It is possible to dispense with the TA and TB timers in this way because the switches are not expected to receive any transmission "gaps" from the failed link channel. The state machines described above still apply however Switches 10_B and 10_D would pass through the WAIT A and WAIT B states instantaneously after detecting any loss of flow on channel C_1. In effect those switches would make a state transition directly from the NORMAL state to a LOCK or SKEW state. The setting of T_{WAIT} to zero is optional: T_{WAIT} may be zero; and T_{WAIT} may be non-zero. Both options are claimed by the invention.

The following provides a detailed description of the delays which are encountered during protection switching. The failure of a link causes a transmission "gap" to be propagated downstream along the link from the point of failure. Gaps with a duration of less than a time T_{WAIT} do not activate protection switching in the switches 10. Only the first downstream switch 10 performs protection switching in response to that loss of flow, and only after a time greater than T_{WAIT}. Where T_{WAIT} is zero protection switching occurs immediately after the loss of flow condition has been

detected. Therefore, when a switch 10 in the NORMAL state detects the loss of flow on either inputs 12 or 16:

1. the switch 10 switches from the NORMAL state into either the WAIT A state or the WAIT B state and as soon as possible thereafter commences transmitting the dummy flow downstream on the failed link so that the duration of the transmission gap sensed downstream is minimised;
2. if the delay time T_{WAIT} is greater than zero, the switch 10 waits one delay time T_{WAIT} to allow any transmission gaps to pass by before switching to one of the LOCK A, LOCK B, SKEW A or SKEW B states;

3a. if any flow is detected again (ie. the flow has been restored for some reason or a dummy flow is detected) the switch 10 reverts immediately to the NORMAL state;

3b. if the flow is still absent at the end of the time T_{WAIT}, the switch 10 switches to one of the LOCK A, LOCK B, SKEW A or SKEW B states.

The delay T_{WAIT} is calculated from the time of the most recent loss of flow. When the flow is intermittent the T_{WAIT} timer is restarted at each instant that the loss of flow is newly detected. When the detected transmission gaps are shorter than T_{WAIT} the switch 10 returns to the NORMAL state at the end of each gap. Only when the flow has been lost for a continuous period greater than T_{WAIT} the switch 10 switches to a SKEW or LOCK state. The delay T_{WAIT} ensures that any switch 10 which may exist downstream from the point of failure does not disrupt unnecessarily the contra-flowing flow. Only the switch 10 adjacent to and immediately downstream from a point of failure switches into the SKEW A or SKEW B state.

Comment	Switch 10_A	Switch 10_B	Flow Received at Site 102	Time Values
Immediately prior to failure	N	N	flow from 100	
After detection of failure at 10_B	N	WB	X	↑ ↑
During switching at 10_B	N	WB	X	T_{SWITCH}
During switching at 10_B	N	WB	X	↓ T_{WAIT}
After switching completed at 10_B	N	WB	dummy flow from 10_B	
During WAIT B state at 10_B	N	WB	dummy flow from 10_B	↓
After Expiry of T_{WAIT} at 10_B	N	SB	dummy flow from 10_B	↑
During propagation delay 43	N	SB	dummy flow from 10_B	propagation delay 43
During propagation delay 43	N	SB	dummy flow from 10_B	↓
After detection of failure at 10_A	WA	SB	dummy flow from 10_B	↑
During and after switching at 10_A	WA	SB	dummy flow from 10_B	T_{WAIT}
During WAIT A state at 10_A	WA	SB	dummy flow from 10_B	↓
After detection of flow at 20 (10_A)	LA	SB	dummy flow from 10_B	↑
During propagation delay 51	LA	SB	dummy flow from 10_B	propagation delay 51
During propagation delay 51	LA	SB	dummy flow from 10_B	↓
After propagation delay 51	LA	LB	X	↑
During switching at 10_B	LA	LB	X	T_{SWITCH}
During switching at 10_B	LA	LB	X	↓
After switching at 10_B	LA	LB	flow from 100	
Steady state	LA	LB	flow from 100	

Table 4: Timing Relationships
Note: No time scale is implied by Table 4.

Figure 18 shows two sites 100 and 102 connected by a network and switches 10_A and 10_B. Table 4 illustrates the delays which occur after the failure of link 41 and the relationship with state transitions at switches 10_A and 10_B. Time in the table passes from the top downwards. "X" represents an indeterminate transitional value which may occur during switching in real devices. Initially both switches 10_A and 10_B are in the NORMAL state and site 100 transmits to site 102 via the link 41. The network is arranged so that the transmission delay between switch 10_B and the receiving site 102 is negligible.

Link 41 fails at location 53. The loss of flow is detected at 53. Switch 10_B takes a finite switching time T_{SWITCH} before it is able to activate and to transmit its dummy flow in state WAIT B towards site 102. If no incoming flow from switch 10_A via link 41 is detected at switch 10_B at the end of time T_{WAIT}, switch 10_B switches to the SKEW B state. After the propagation delay incurred in traversing link 43 the flow at input 16 of switch 10_A is lost. Switch 10_A switches from the NORMAL state to the WAIT A state. Let switch 10_A detect at its input 20 the incoming flow from switch 10_B while still in the WAIT A state. Switch 10_A switches to the LOCK A state at the end of the delay T_{WAIT}. The flow on link 49 presented to input 12 of switch 10_A is then transmitted to switch 10_B via the protection link 51. When detected at input 20 of switch 10_B after the propagation delay incurred in traversing link 51, switch 10_B switches from the SKEW B state to the LOCK B state, completing the repair action. The required flow from site 100 is again received at site 102.

Figure 19 illustrates the relationship between the time delays. At the top is shown a digital link flow whose value alternates between a high value 1 and a low value 0 as time passes to the right. The illustration is not to scale. The illustration may represent a digital optical flow which turns ON and OFF, or an electrical flow which alternates between a high voltage and a low voltage. The time 104 is the longest expected zero sequence in a normal flow. To successfully detect the loss of flow condition the switch 10 must fail to detect a flow for at least the time T_{DETECT} 106, which consists of the longest expected zero sequence 104 together with an error tolerance 108. The longest expected zero sequence 104 is the longest expected period for which received flow may be absent from a healthy flow including a dummy flow. Transmission gaps with a duration of less than a time T_{DETECT} 106 are not interpreted as “linked failed”. The absence of a flow for greater than T_{DETECT} 106 is interpreted as "link failed". Therefore T_{DETECT} is dependent upon the transmission bit rate and the line coding scheme used. For example, in a digital optical embodiment in which logical "zeros" are represented by the absence of light, T_{DETECT} would be greater than the duration of the greatest expected number of consecutive optical "zeros" divided by the line rate measured in bits per unit time at the lowest transmission bit rate for which the module is designed.

The figure shows a healthy flow. If on the other hand the flow remained at zero for at least the time T_{DETECT} 106 then the flow would be deemed lost and one or more switch 10 state changes would be initiated. At the instant marked 110 at the end of T_{DETECT} 106 if the flow is still absent (not shown in the figure) a switch 10 switches from the NORMAL state to a WAIT state.

Any realisation of the switch 10 requires a finite time to perform switching operations and to activate and transmit the dummy flow if necessary. T_{SWITCH} 112 is the maximum time required to implement any changes of state at any switch 10 in a network. T_{SWITCH} 112 is equal to or greater than the time required to activate, switch and transmit the dummy flow and to perform all other switching functions. The longest possible transmission gap which is propagated downstream is less than or equal to the sum of T_{DETECT} 106 and T_{SWITCH} 112. Downstream switches 10 remain in a WAIT state for at least a time T_{WAIT} in order to allow any transmission gaps from upstream to flow by. In any realisation of a switch 10 with T_{WAIT} not equal to zero, T_{WAIT} 114 is therefore required to be equal to or greater than the sum of: i) the switching time T_{SWITCH} 112; ii) T_{DETECT} 106; and iii) the further delay 116, which provides a sufficiently large margin so that T_{WAIT} at any switch 10 in the network is always greater than the longest possible transmission gap.

For example: if the longest expected zero sequence 104 on a 10Km 155·52 Mbps optical point to point link is four optical zeros then T_{DETECT} 106 would be greater than the time 104 and would be given by:

$$T_{DETECT} \quad > \quad 4 \ / \ (\ 155{\cdot}52 \times 10^6\ ) \text{ seconds} \quad = \quad 25{\cdot}7 \text{ nanoseconds}$$

With a 74 nanosecond error tolerance 108 the decision to switch to the WAIT A or WAIT B state from the NORMAL state may be made within 100 nanoseconds of the detection of loss of flow. With an additional maximum switching time T_{SWITCH} 112 also of 100 nanoseconds the maximum duration of any transmission gap would be 200 nanoseconds. Setting the TA and TB timer delays T_{WAIT} 114 to 500 nanoseconds each provides a 300 nanosecond margin 116. Switching to a SKEW or LOCK state

would occur within 500 nanoseconds of the detection of loss of flow at each Network Repair Module. If the flow travelled at 2×10^8 m/sec then the round trip flow propagation delay would be

$$2 \times 10^4 \ / \ 2 \times 10^8 \text{ seconds} = 100 \ \mu\text{seconds}$$

Therefore in the worst case which occurs when one link fails near a transmitter protection switching would be completed within 101 μseconds. In the best case which occurs when each link fails near a receiver protection switching would occur within 500 nanoseconds. In the best case the protected flow in one direction would be switched to and transmitted along the redundant link almost immediately after the link failure. The receiving site downstream would receive the flow from the redundant link/channel almost immediately after detecting the loss of flow (assuming that the transmission delays of the protected link and the redundant link were comparable). By making the redundant link shorter than the protected link, the receiver would receive the diverted flow from the redundant link before the receiver lost the flow from the normal link. Setting the TA and TB timer delays T_{WAIT} to zero furnishes very high speed protection switching. For the example above, protection switching would be completed in less than 200 nanoseconds at each switch 10. Better performance would be achieved by reducing the generous margins given in this example. When the propagation delay in the protected link is identical to the propagation delay in the redundant link it is possible to switch paths without losing bit or frame synchronisation at the receiving sites.

After the failed links and sites have been repaired flows would be restored to the links which were originally used. Restoration of flows is achieved by forcing one or more switches 10 into the NORMAL state. A forced state transition at one switch 10 may automatically initiate state transitions at other remote switches 10. The forced switch 10 together with other switches 10 co-operatively achieve restoration. It is not necessary to force state transitions at all switches 10. For example the restoration of a flow to one of the pair of links 37 from one of the redundant links of channel C_2 in Figure 17 after the failure and subsequent repair of channel C_1 would

be achieved by forcing only one local switch 10_D at for example the exchange end into the NORMAL state. The loss of flow from the redundant links and the detection again of a flow on the normal link at the remote switch 10_B together are sufficient to cause switch 10_B also to switch to the NORMAL state, so completing restoration.

Switches 10 and switching systems in accordance with this invention can thus be controlled from remote locations. This feature is of benefit to network operators who would be spared the expenses of additional network management networks, and travelling of technicians to switches 10 at remote sites.

Networks would be protected against the failure of switches 10 either: by installing switches together with separate active or passive by-pass devices; or by incorporating a passive by-pass device into switches. Figure 20 shows a separate by-pass device 118. Operation of the by-pass device would be controlled either: directly by the switch 10_A via a control line 120 as shown; or by some other device. Failure of the switch 10_A or its power supply would cause the by-pass device 118 to be switched so that the bidirectional through paths 122 and 124 are enabled and so that the failed switch 10_A is by-passed. When in the by-passed state any flow or flows received at input 126 are transferred to output 128 and any flows received at input 130 are transferred to output 132. When not in the by-passed state the flows at inputs 126 and 130 are transferred to the switch 10_A. The failure of any power supply which may be required by the by-pass device 118 would also cause the switch 10A to be by-passed.

Now that embodiments of this invention have been described in detail it will be apparent to those skilled in the relevant arts that numerous modifications and variations may be made without departing from the basic inventive concepts. For example, embodiments of the invention may be realised using a microprocessor to control the switches 10 instead of a field programmable gate array. Indeed, embodiments of the invention may be realised entirely in software that may be hosted in a computer, communications equipment or other equipment that may be

installed at sites in a network. The software would implement the invention by altering routing tables, configuration tables or other mapping tables which specify associations between incoming and outgoing links and any of packet source addresses; packet destination addresses; telephone numbers; sockets; ports; incoming links; outgoing links; node addresses; connections; virtual channels; virtual paths; SDH tributary units (TU-ns); SONET virtual tributaries (VT-ns); user group multiplex labels; channel identifiers; routes; and other types of channels. Channels, paths, connections and routes would thus be remapped to links that connect to healthy parts of the network in response to the detection of a loss of flow condition. The loss of flow condition may be detected either by software or by signals from line interfaces and other devices within the host.

It is also important to note that embodiments of the invention may be structured in a manner so that say P_1 and P_3 would support links with different characteristics from the links at the other port P_2. For example one embodiment of the switch 10 for telecommunications networks would include high power interfaces at P_1 and P_3 that support long range flows, while the remaining port P_2 will include a low power interface for short range flows. This arrangement would enable the switch 10 to be installed in an exchange building or central office between existing transmission equipment and the transmission line. The low power port P_2 would connect to nearby equipment over shorter range links. The high powered ports P_1, P_3 will connect to the longer range protected and protection lengths. Further, in embodiments of the switch 10 line interfaces can be used at the ports in which the band width transmitted onto one or more outgoing links is not equal to the band width received from an incoming link. Switches 10 may thus be constructed for installation in networks that support asymmetric digital subscribed lines (ADSLs).

Also, while embodiments of the invention have been described incorporating a dummy flow generator such a generator is not essential for the protection of point to point links and other network topologies.

Where a synchronising timing structure is embedded within a flow embodiments of the invention may include a time extraction mechanism that recovers synchronising timing from the flow. The timings though recovered may be employed to synchronise the transmission of one or more outgoing flows from a switch. The use of timing recovery and flow generation would enable the maximum number of switches 10 that may be concatenated along a set of links to be greater than embodiments that do not perform regeneration. A wide range of timing recovery components is available both as discrete components or integrated into receiver components.

Embodiments of the switch may include a device or system that monitors and optionally performs switching as a result of fault status or alarm signals furnished by and received from other sites in the network. For example, a switch 10 may be constructed that performs switching functions in response to the receipt of fault status, alarm or other signals.

Embodiments of the switch 10 can also be arranged or operated to detect the restoration of a normally used link after having failed, and automatically re-route a flow back onto that normal link. The requirement for manual intervention to divert a flow from a protection link to a protected link would thus be eliminated. The detection of link restoration may be achieved by transmitting a signal onto a failed link which would be received after restoration of that link.

One range of embodiments of the switch 10 may be constructed to protect the flow of the electrical energy through transmission and distribution networks that span or are contained within one or more continents or nations, a smaller region, a vehicle or machine, a circuit board, and a silicon or other microchip. Each switch 10 would generate a flow of electrical protection signals and transmit onto the electrical power transmission conductors, in both directions. Each physical link in the transmission network would facilitate both a flow of electrical energy in either or both directions; and, a flow of protection signals in either or both directions. The switches 10 would divert the protected power flow onto healthy links in response to lose of receipt of

protection signals.

Finally, as mentioned above, the switches 10 can be constructed to protect networks other than electrical/electro-magnetic communication systems for example to protect networks of vacuum tubes that transport pneumatically propelled canisters; to protect blood vessels and other carriageways in human and other bodies to retain full connectivity between say the heart or other pump and other parts of the body. In this embodiment, the detection means would detect characteristics of blood flow such as blood pressure. Loss of blood pressure along one blood vessel would activate bypass switching to a standby redundant path/blood vessel. A completely distributed network of blood carrying tubes or conduits connected via switches 10 could be constructed to create an extremely robust being.

All such modifications and variations together with others that would be obvious to a person of ordinary skill in the art are deemed to be within the scope of the present invention the nature of which is to be determined from the above description.

Dated this 25th day of August 1999

NILS MARCHANT
By His Patent Attorneys

GRIFFITH HACK
Fellows Institute of Patent and Trade Mark
Attorneys of Australia

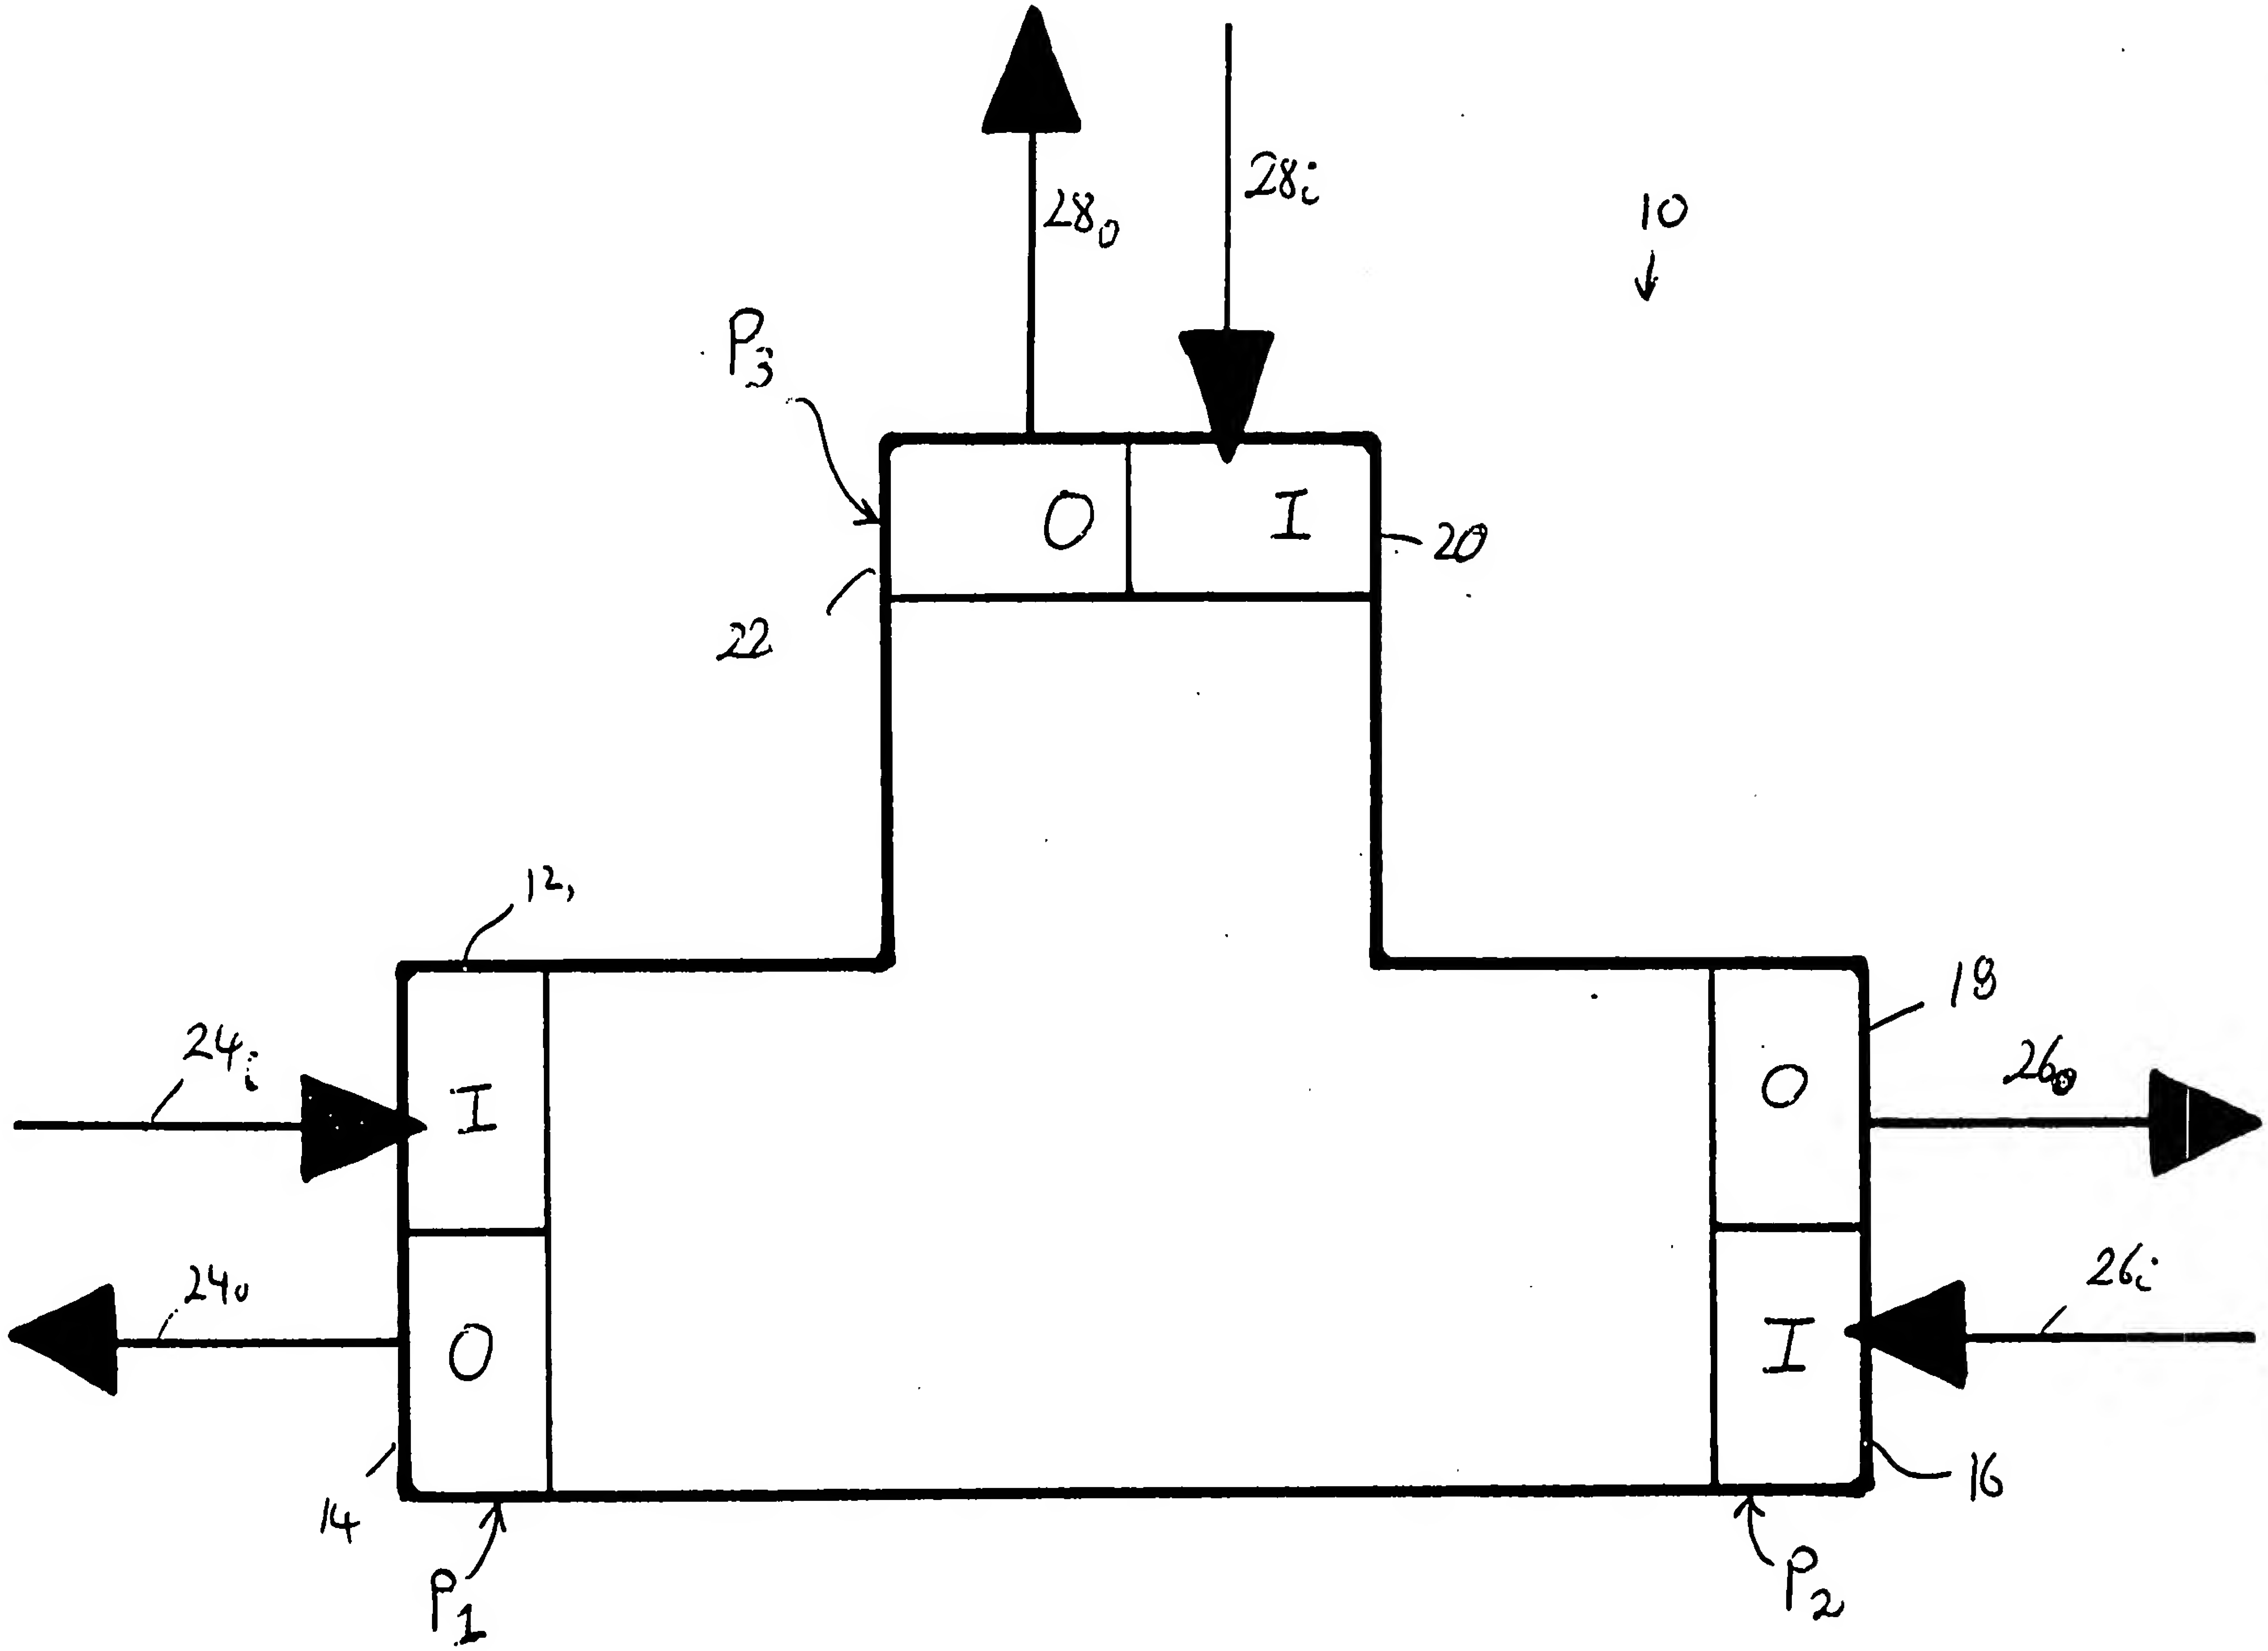

FIG 1

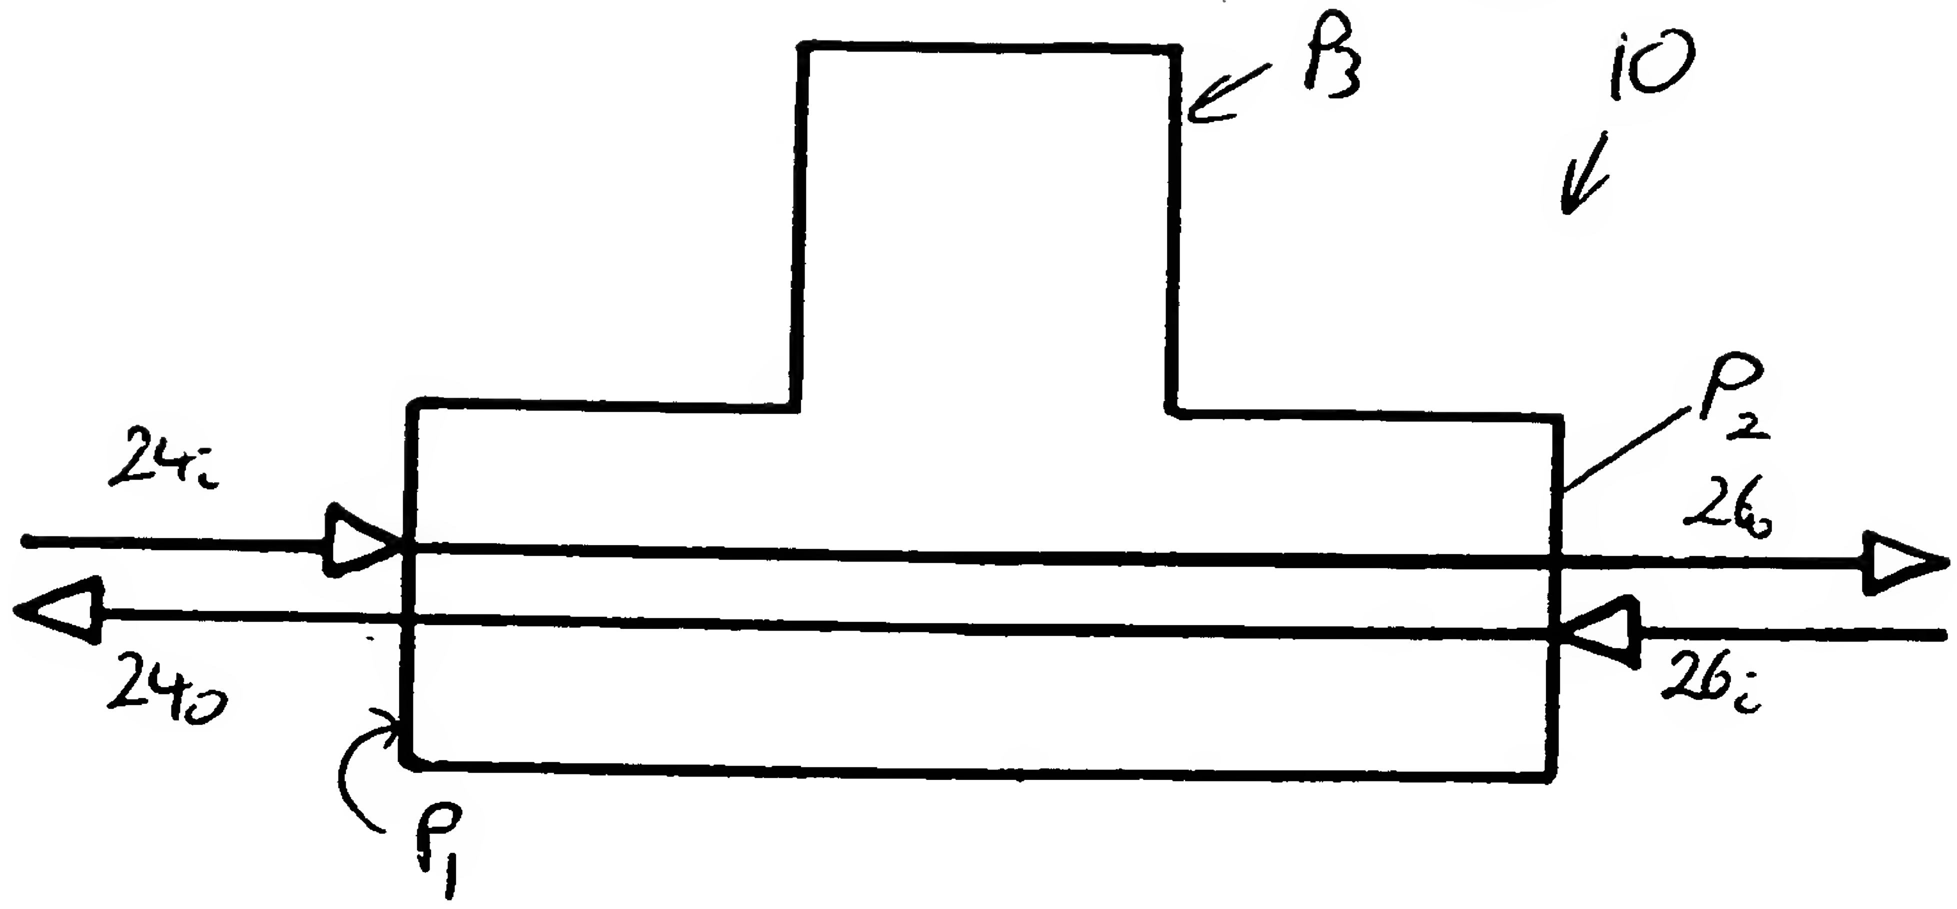

FIG 2A

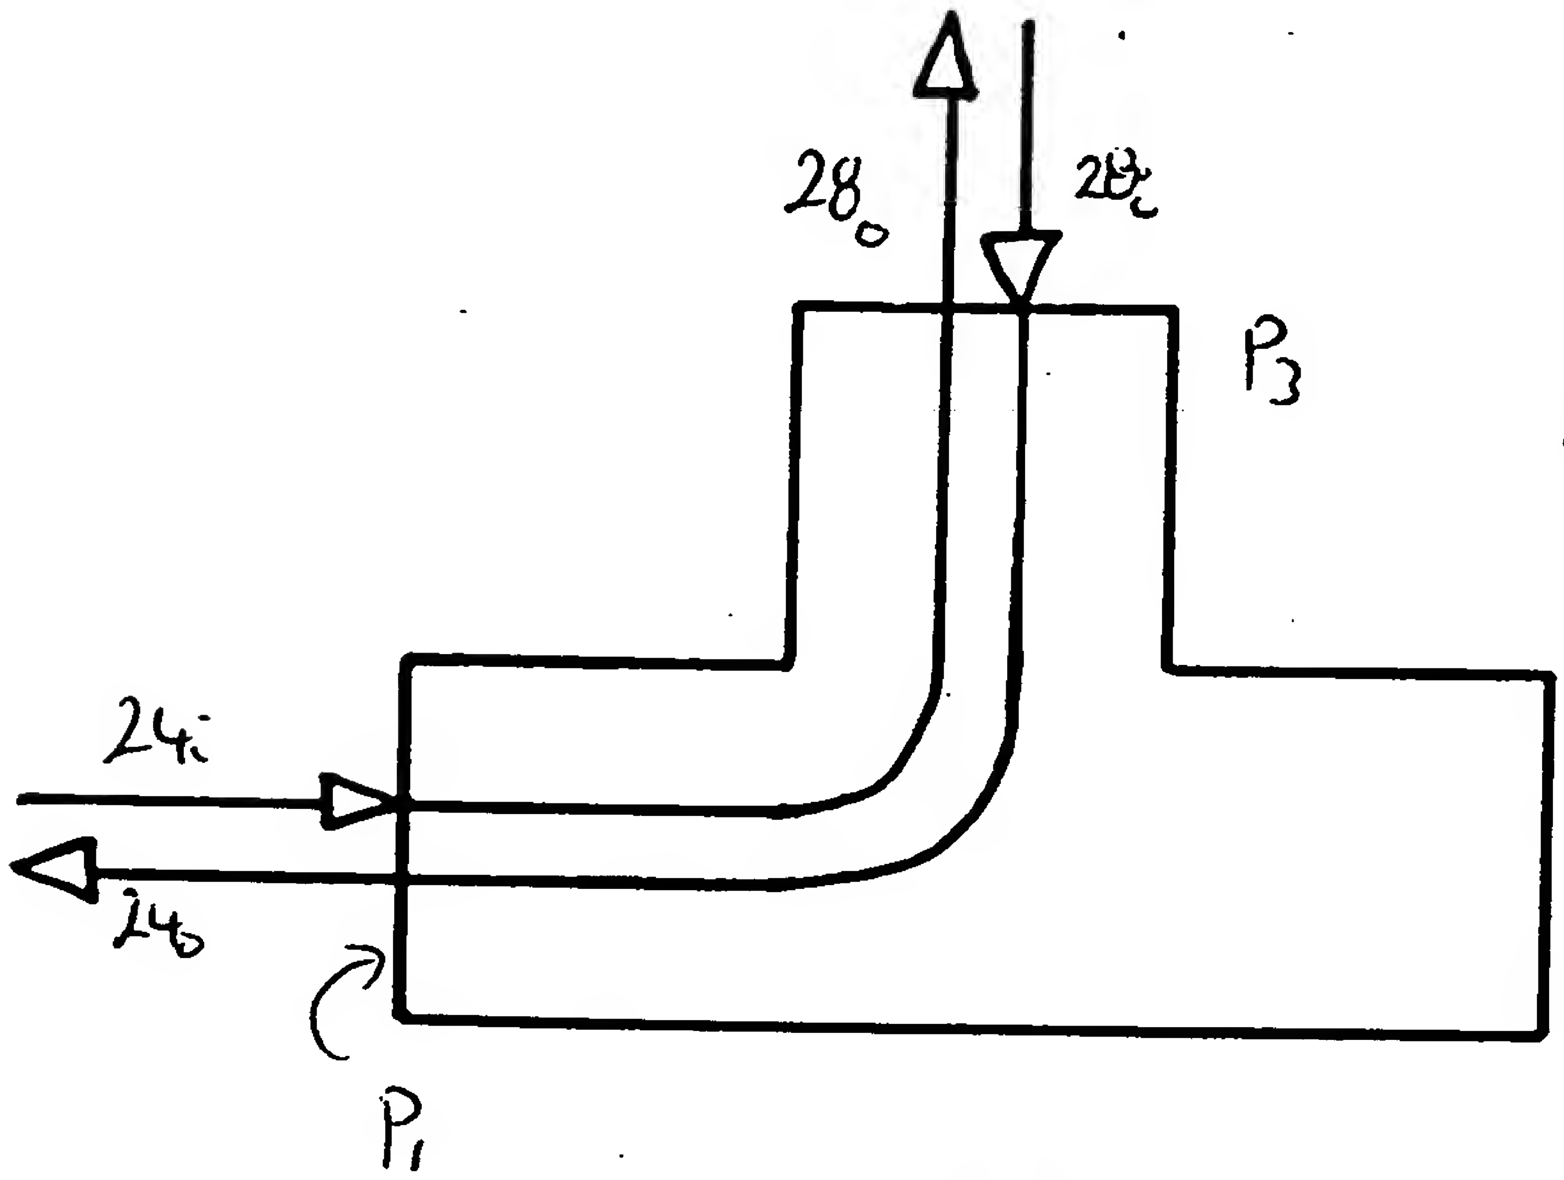

FIG 2B

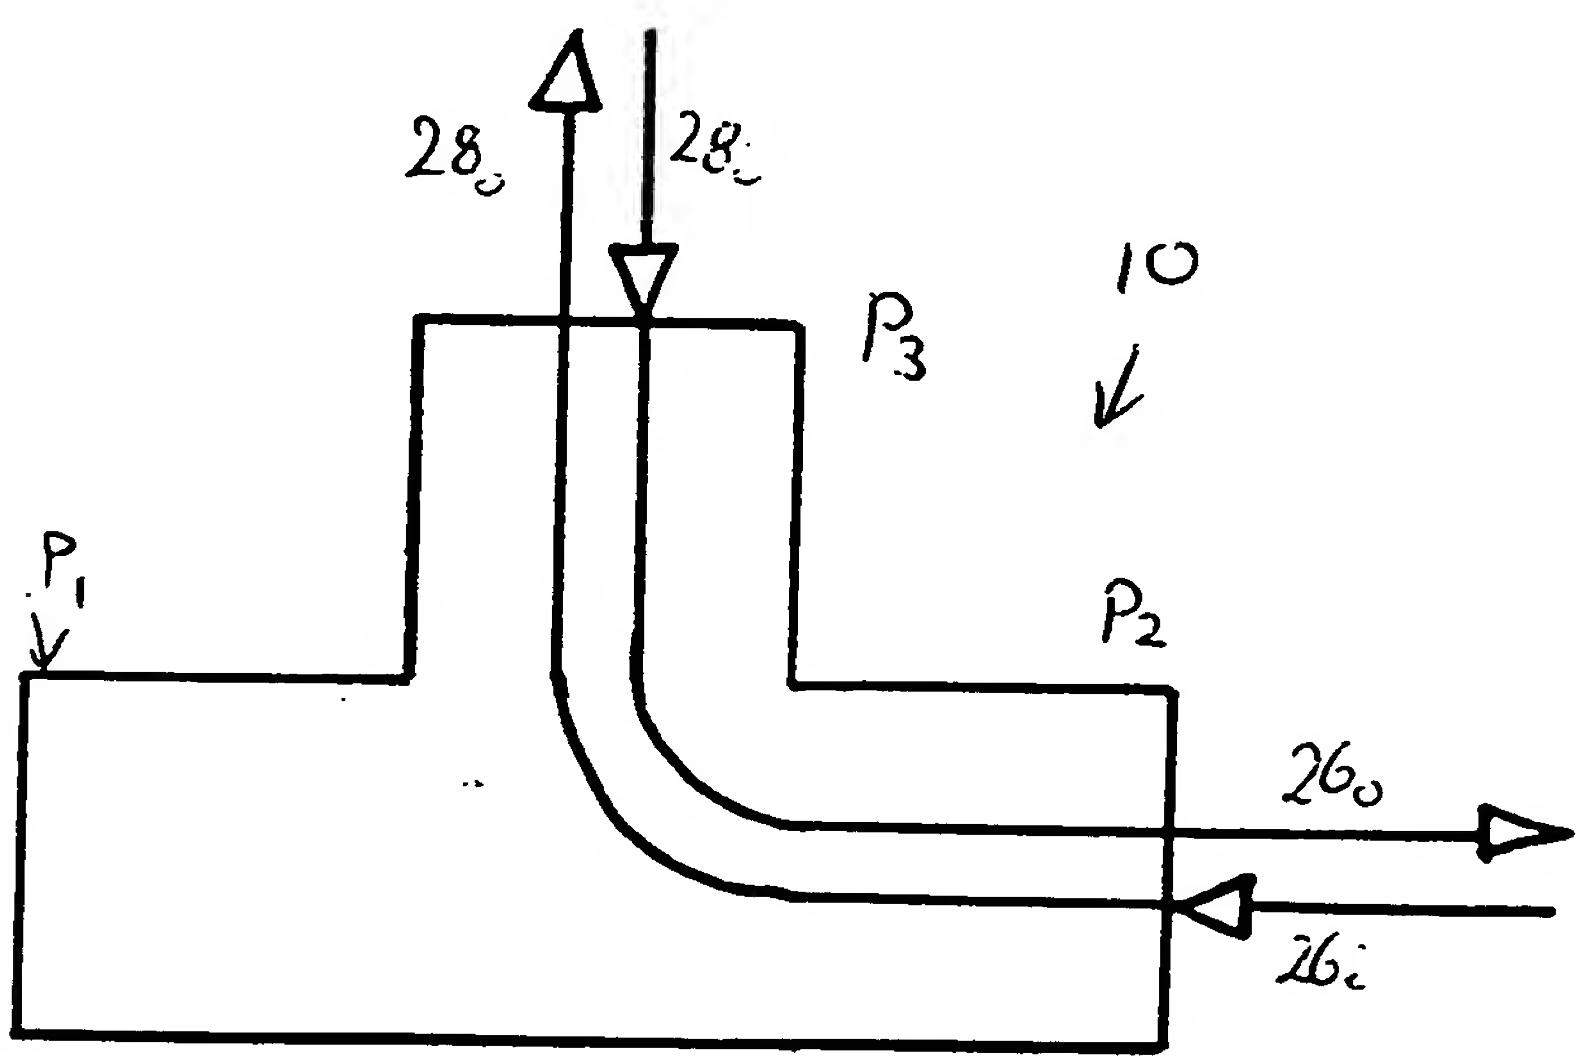

FIG 2C

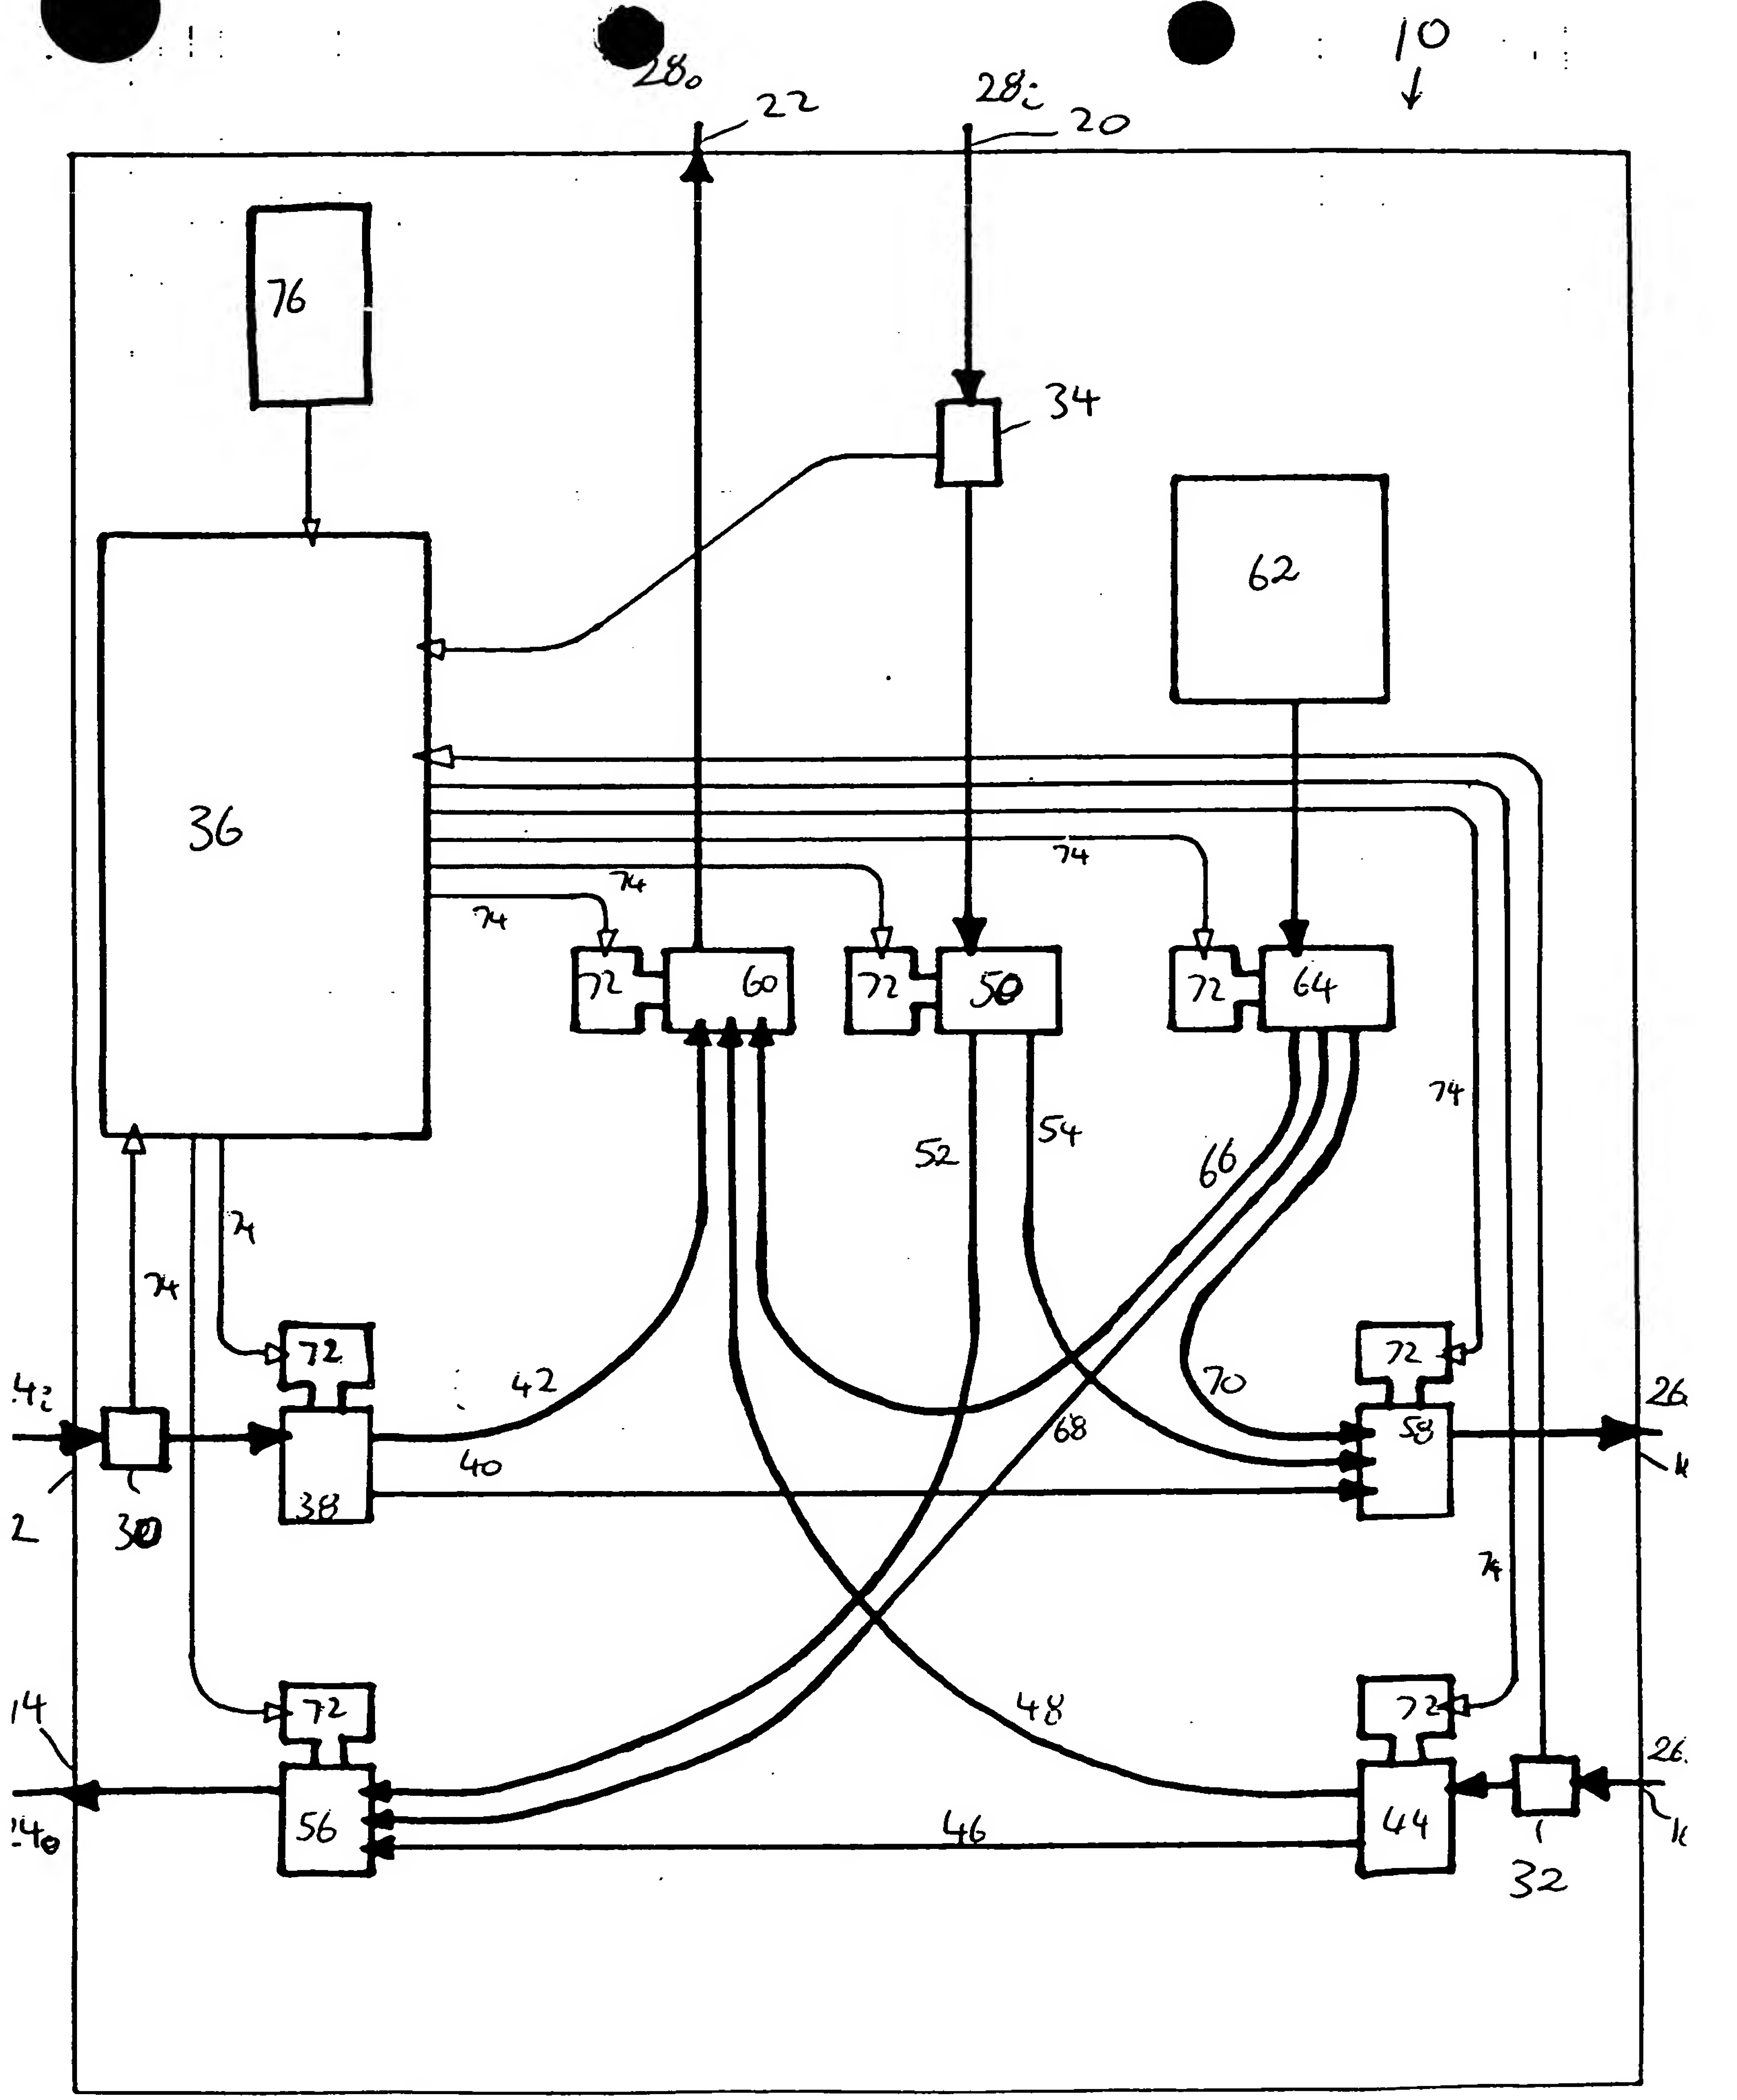
10
28o
22
28i
20
76
34
62
36
74
72
60
50
64
52
54
66
42
40
38
30
70
58
68
48
56
44
46
32
14

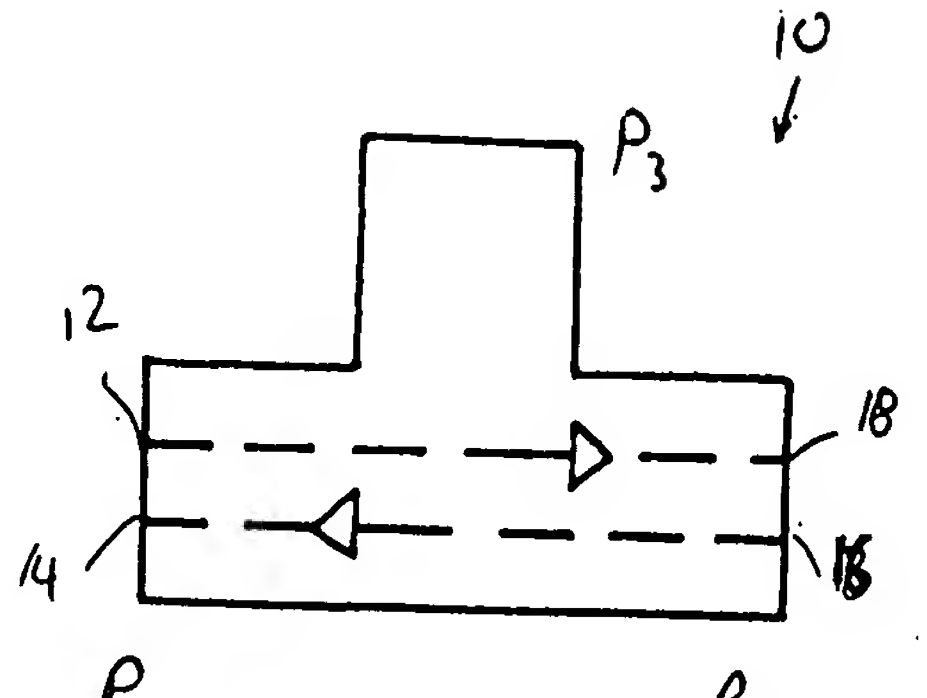

FIG 4

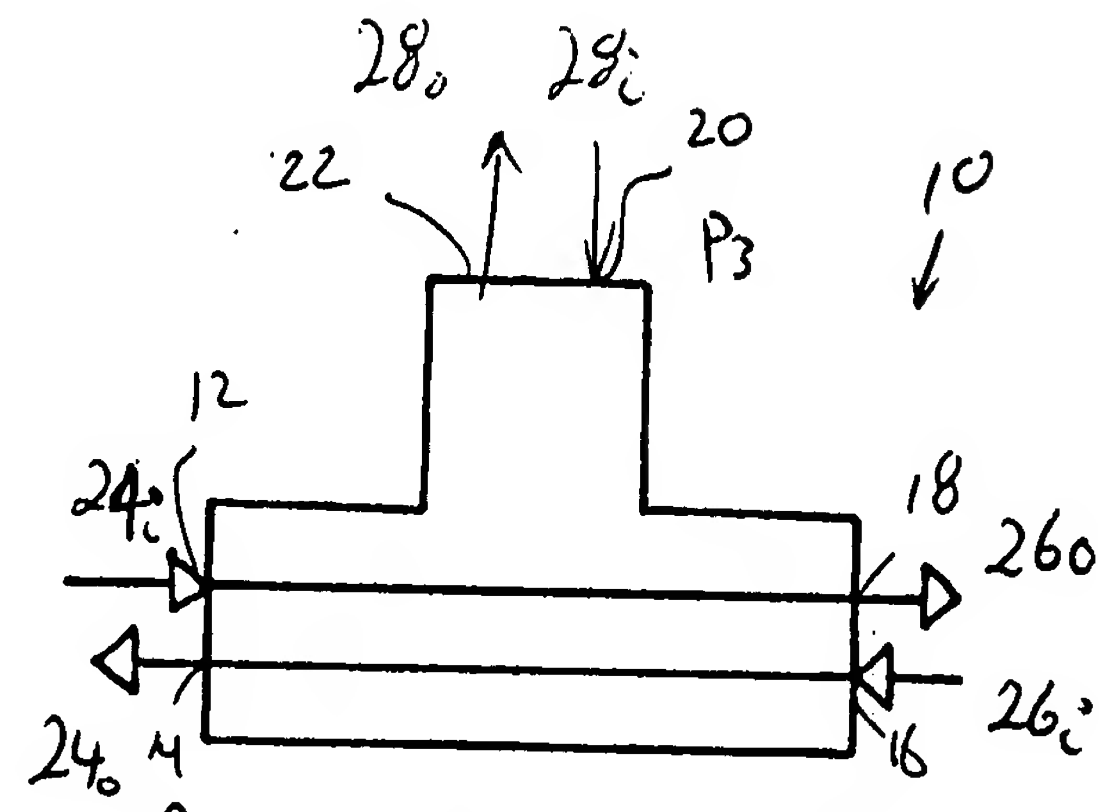

FIG 5

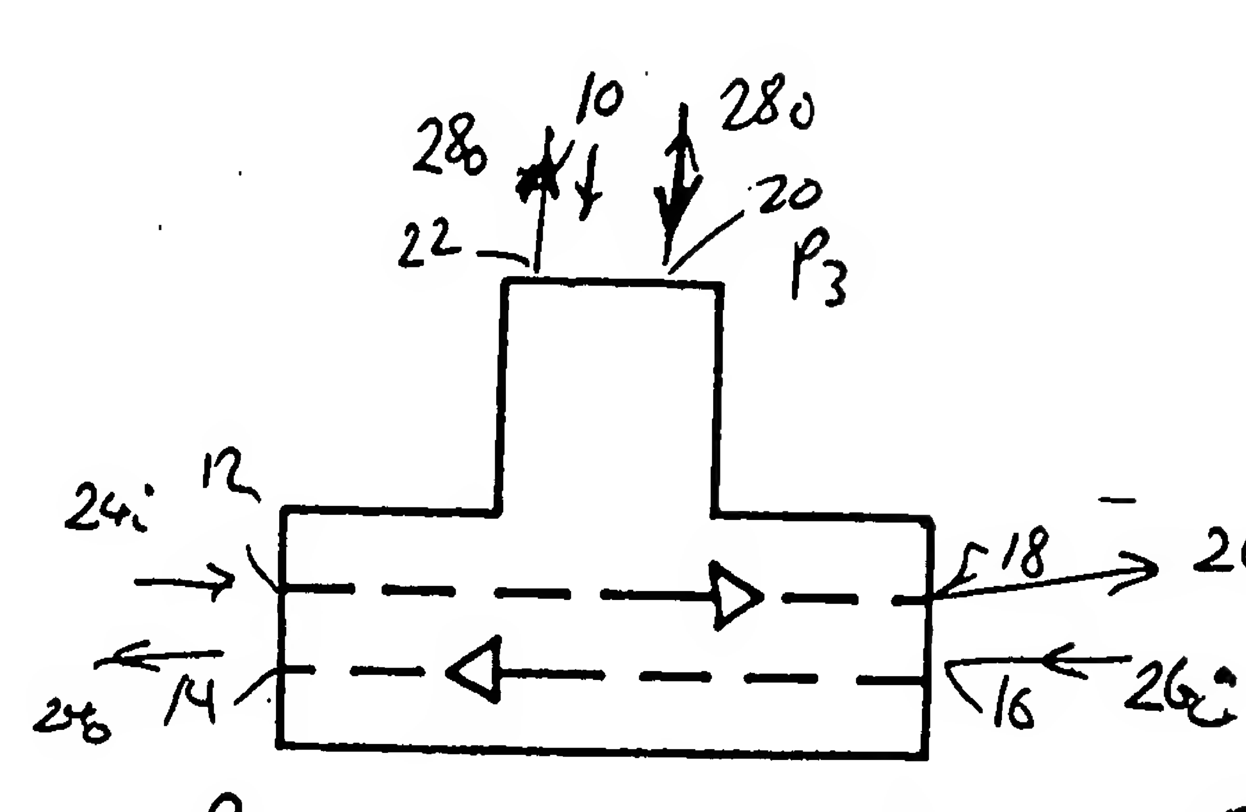

FIG 6

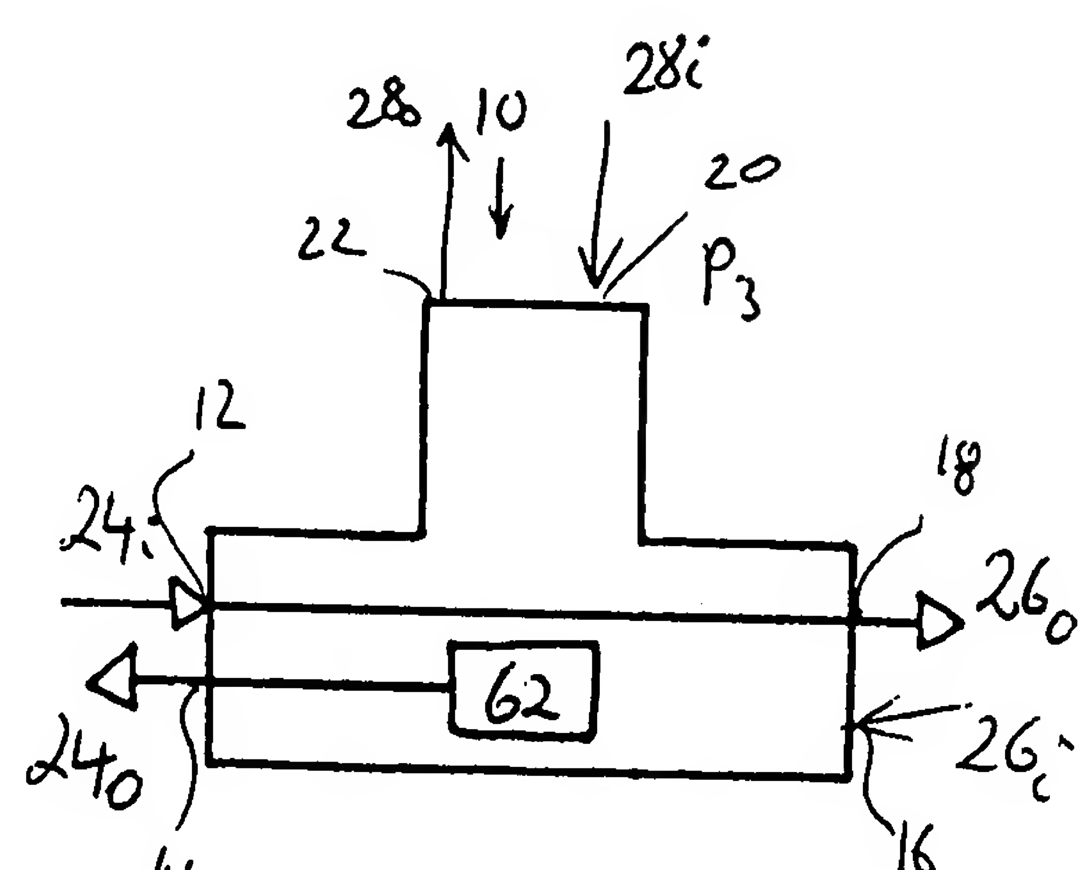

FIG 7

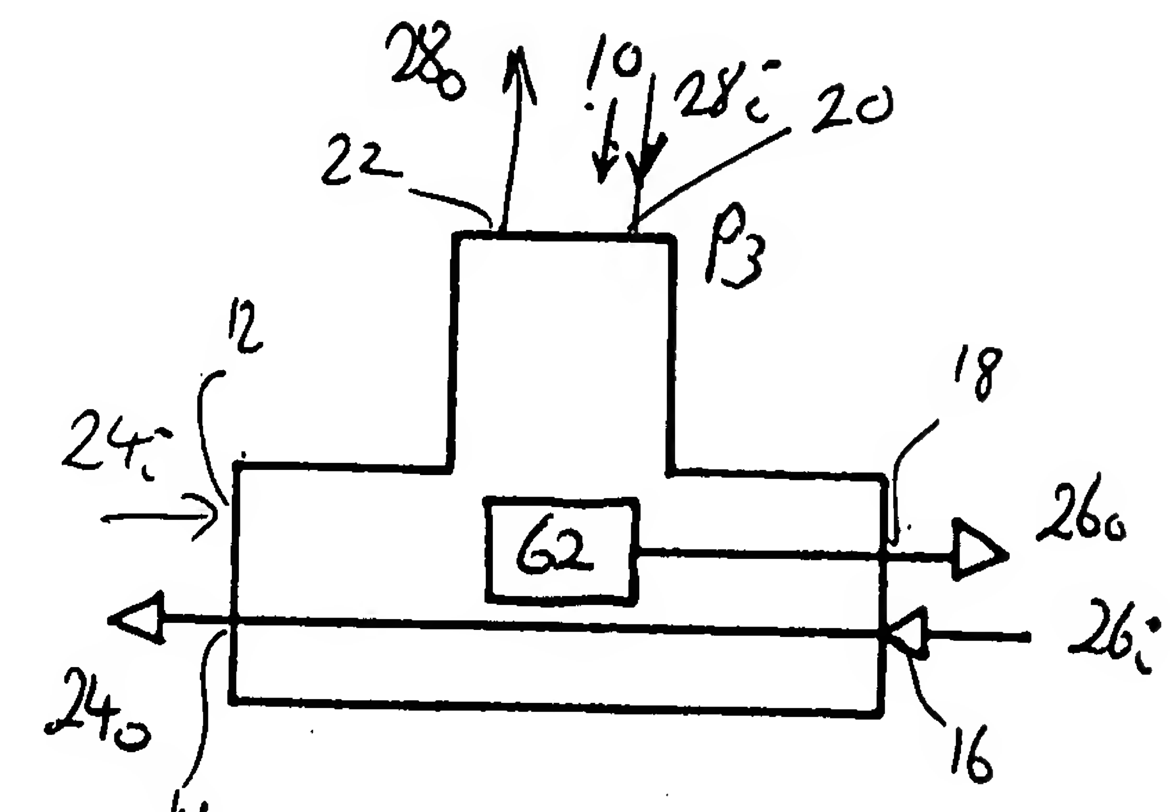

FIG 8

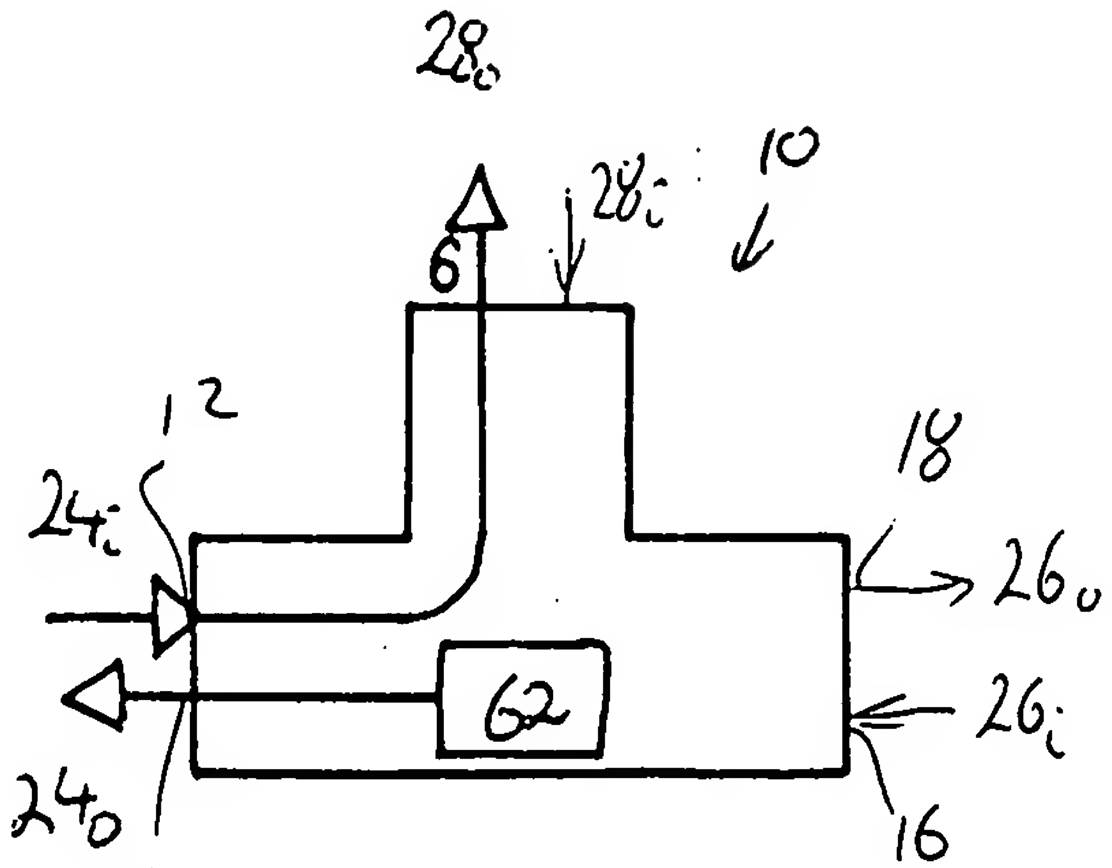

FIG 9

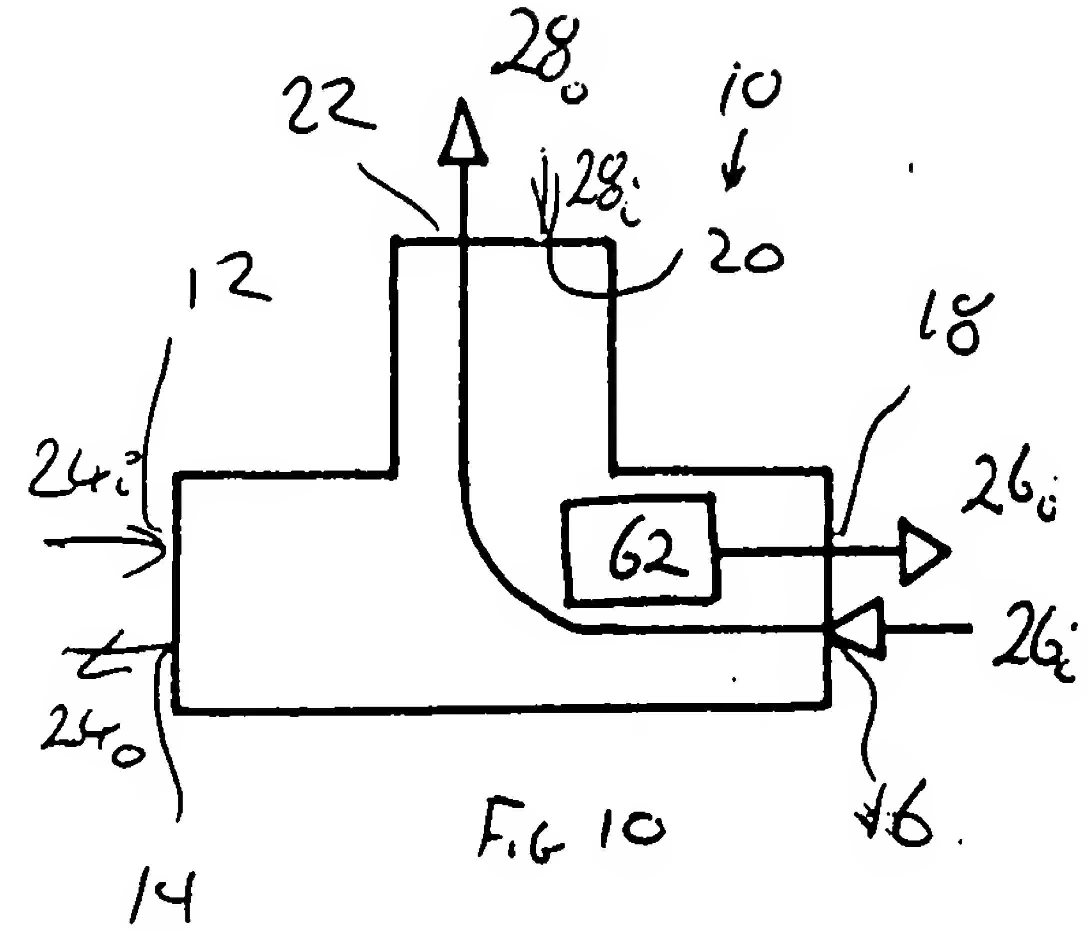

FIG 10

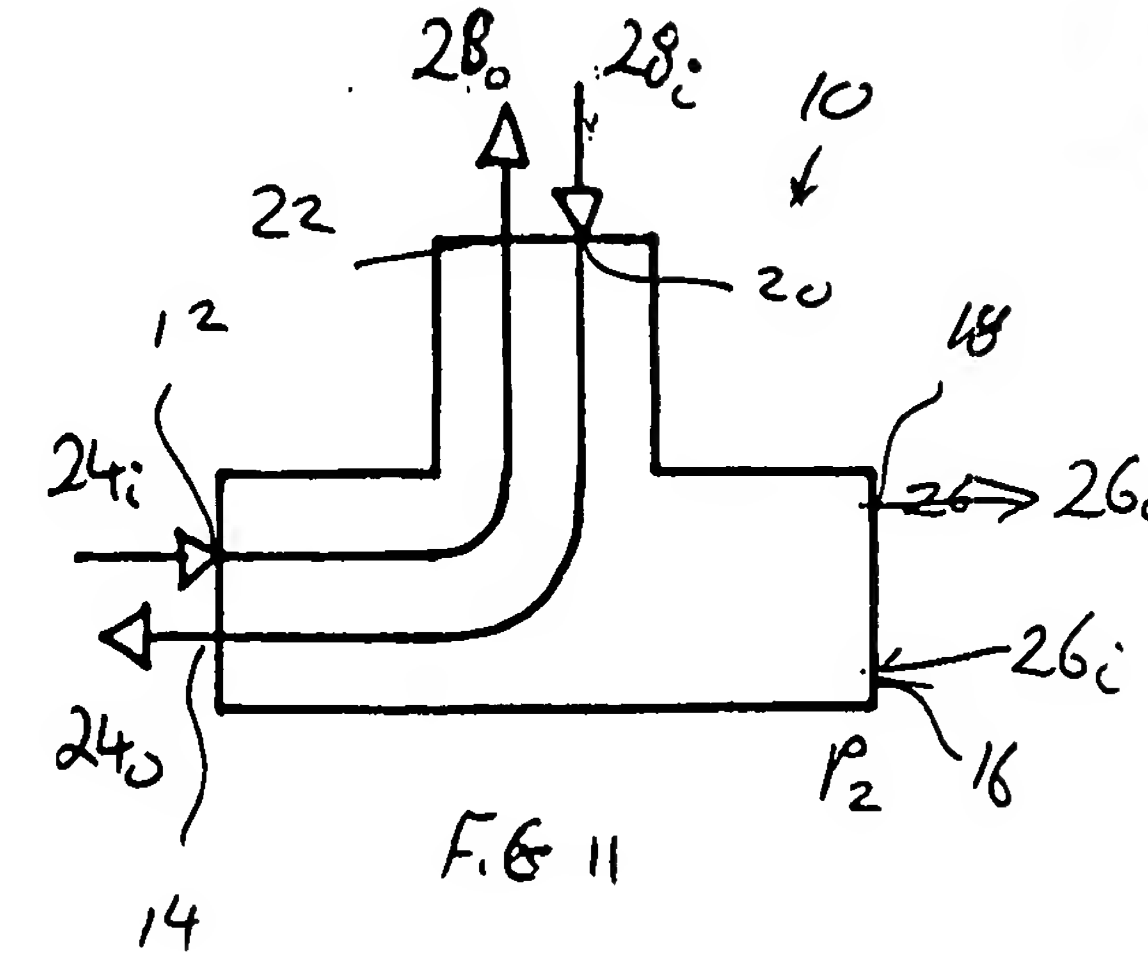

FIG 11

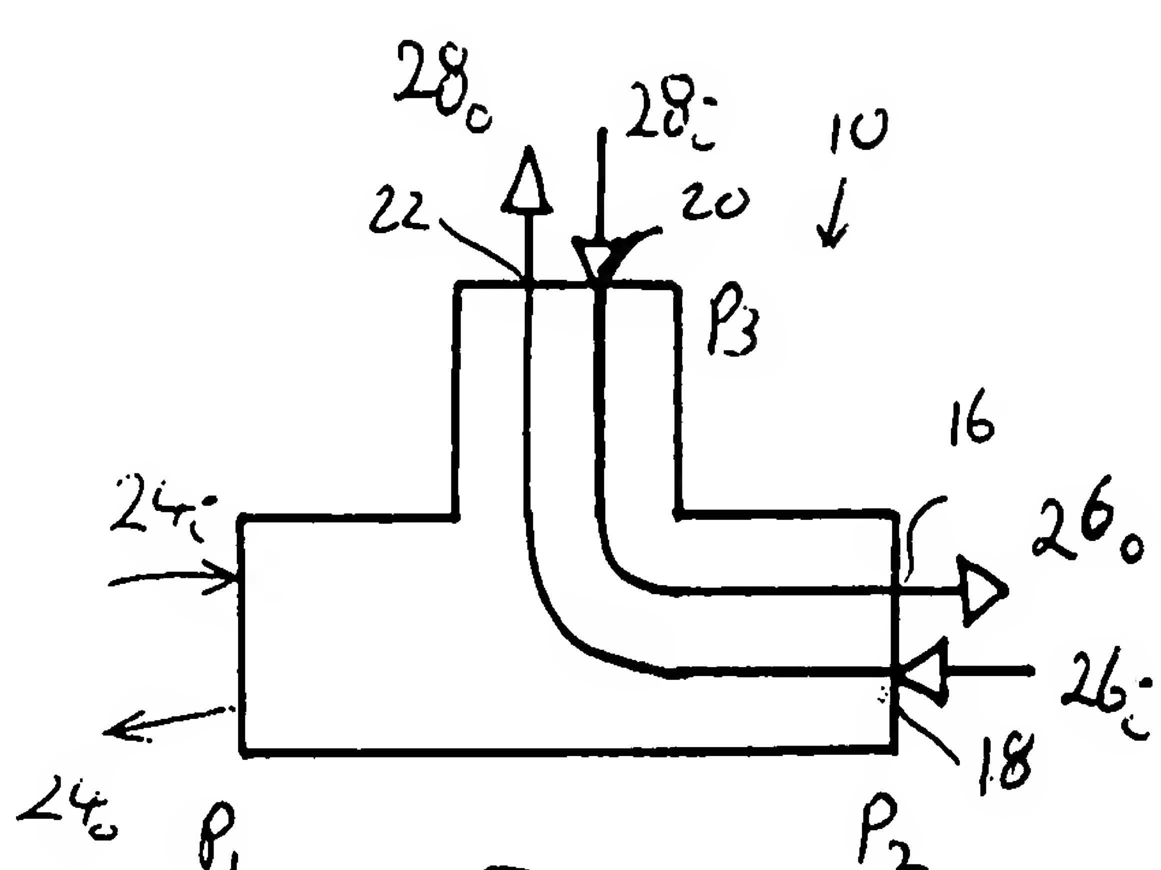

FIG 12

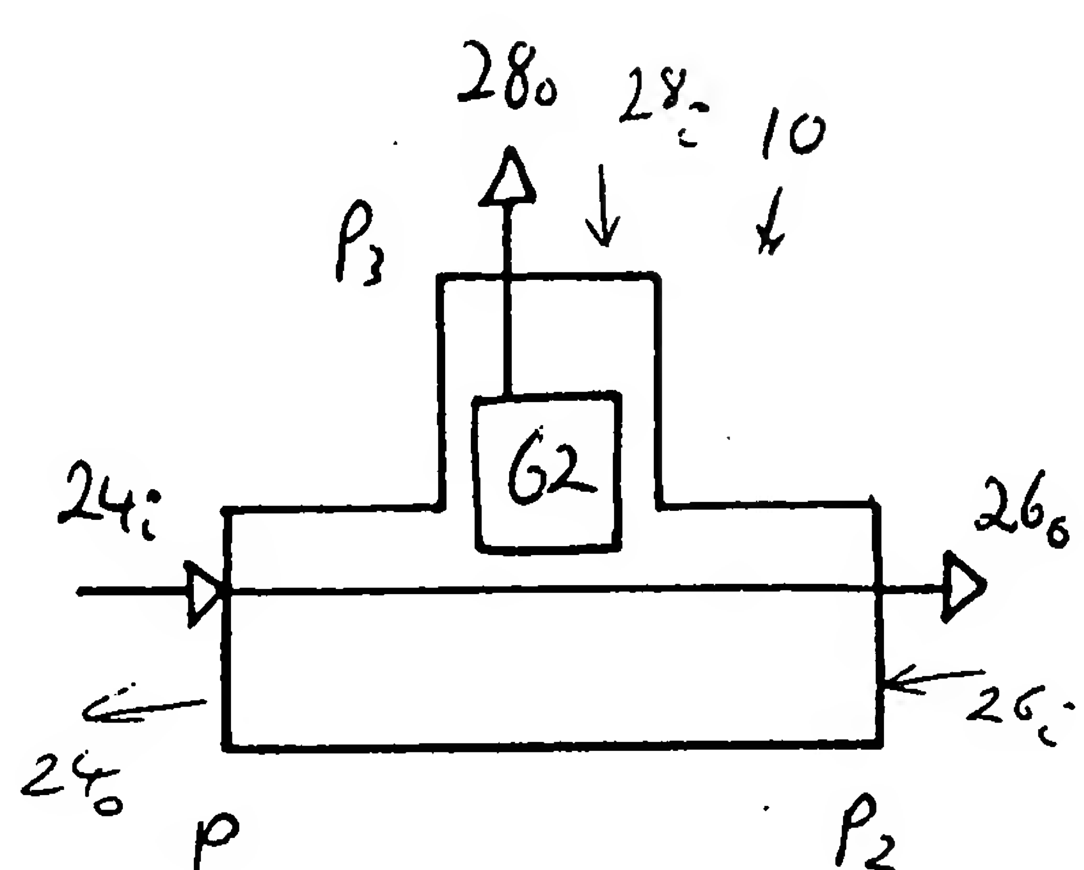

FIG 13

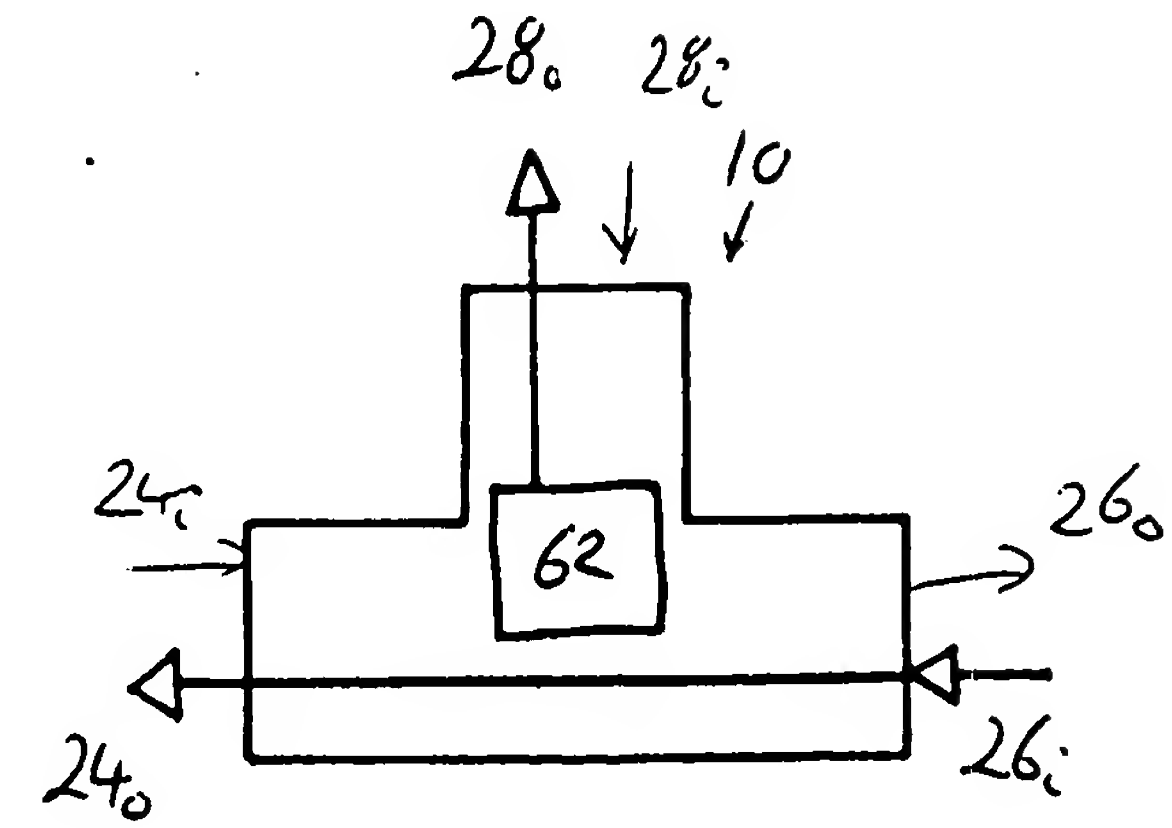

FIG 14

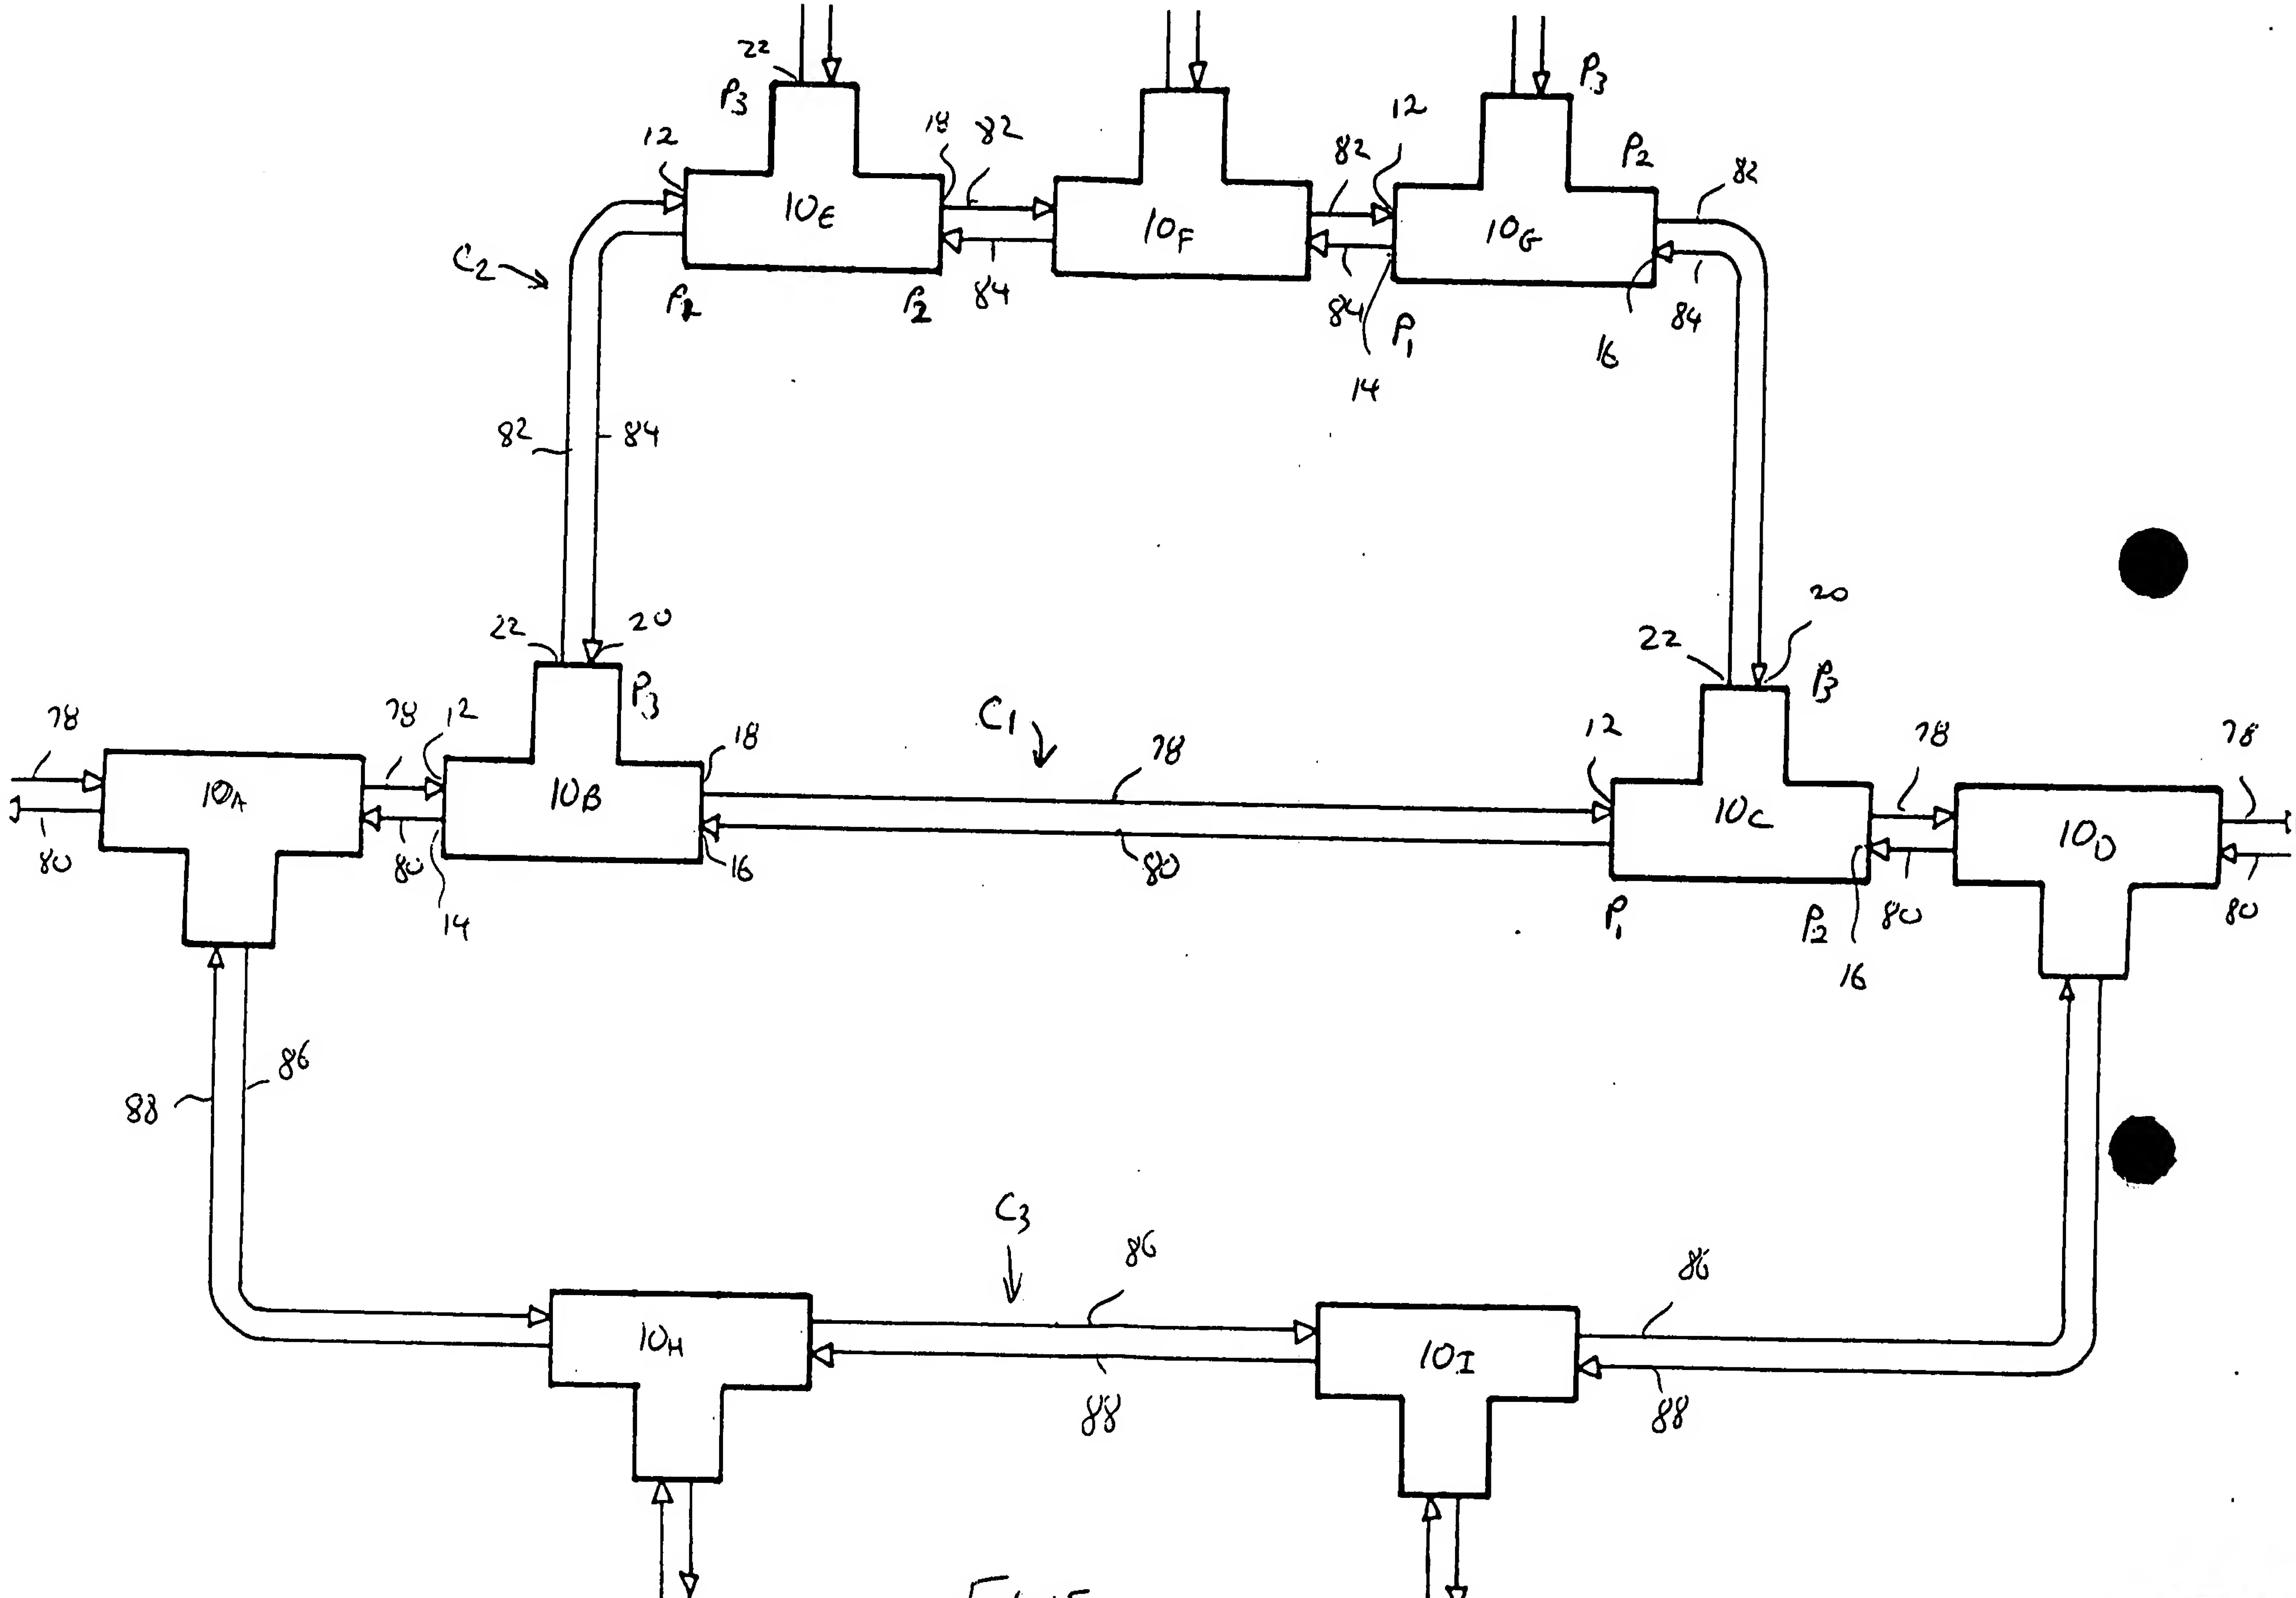

10E
10F
10G
10A
10B
10C
10D
10H
10I
C1
C2
C3
P1
P2
P3
12
14
16
18
20
22
78
80
82
84
86
88

FIG 15

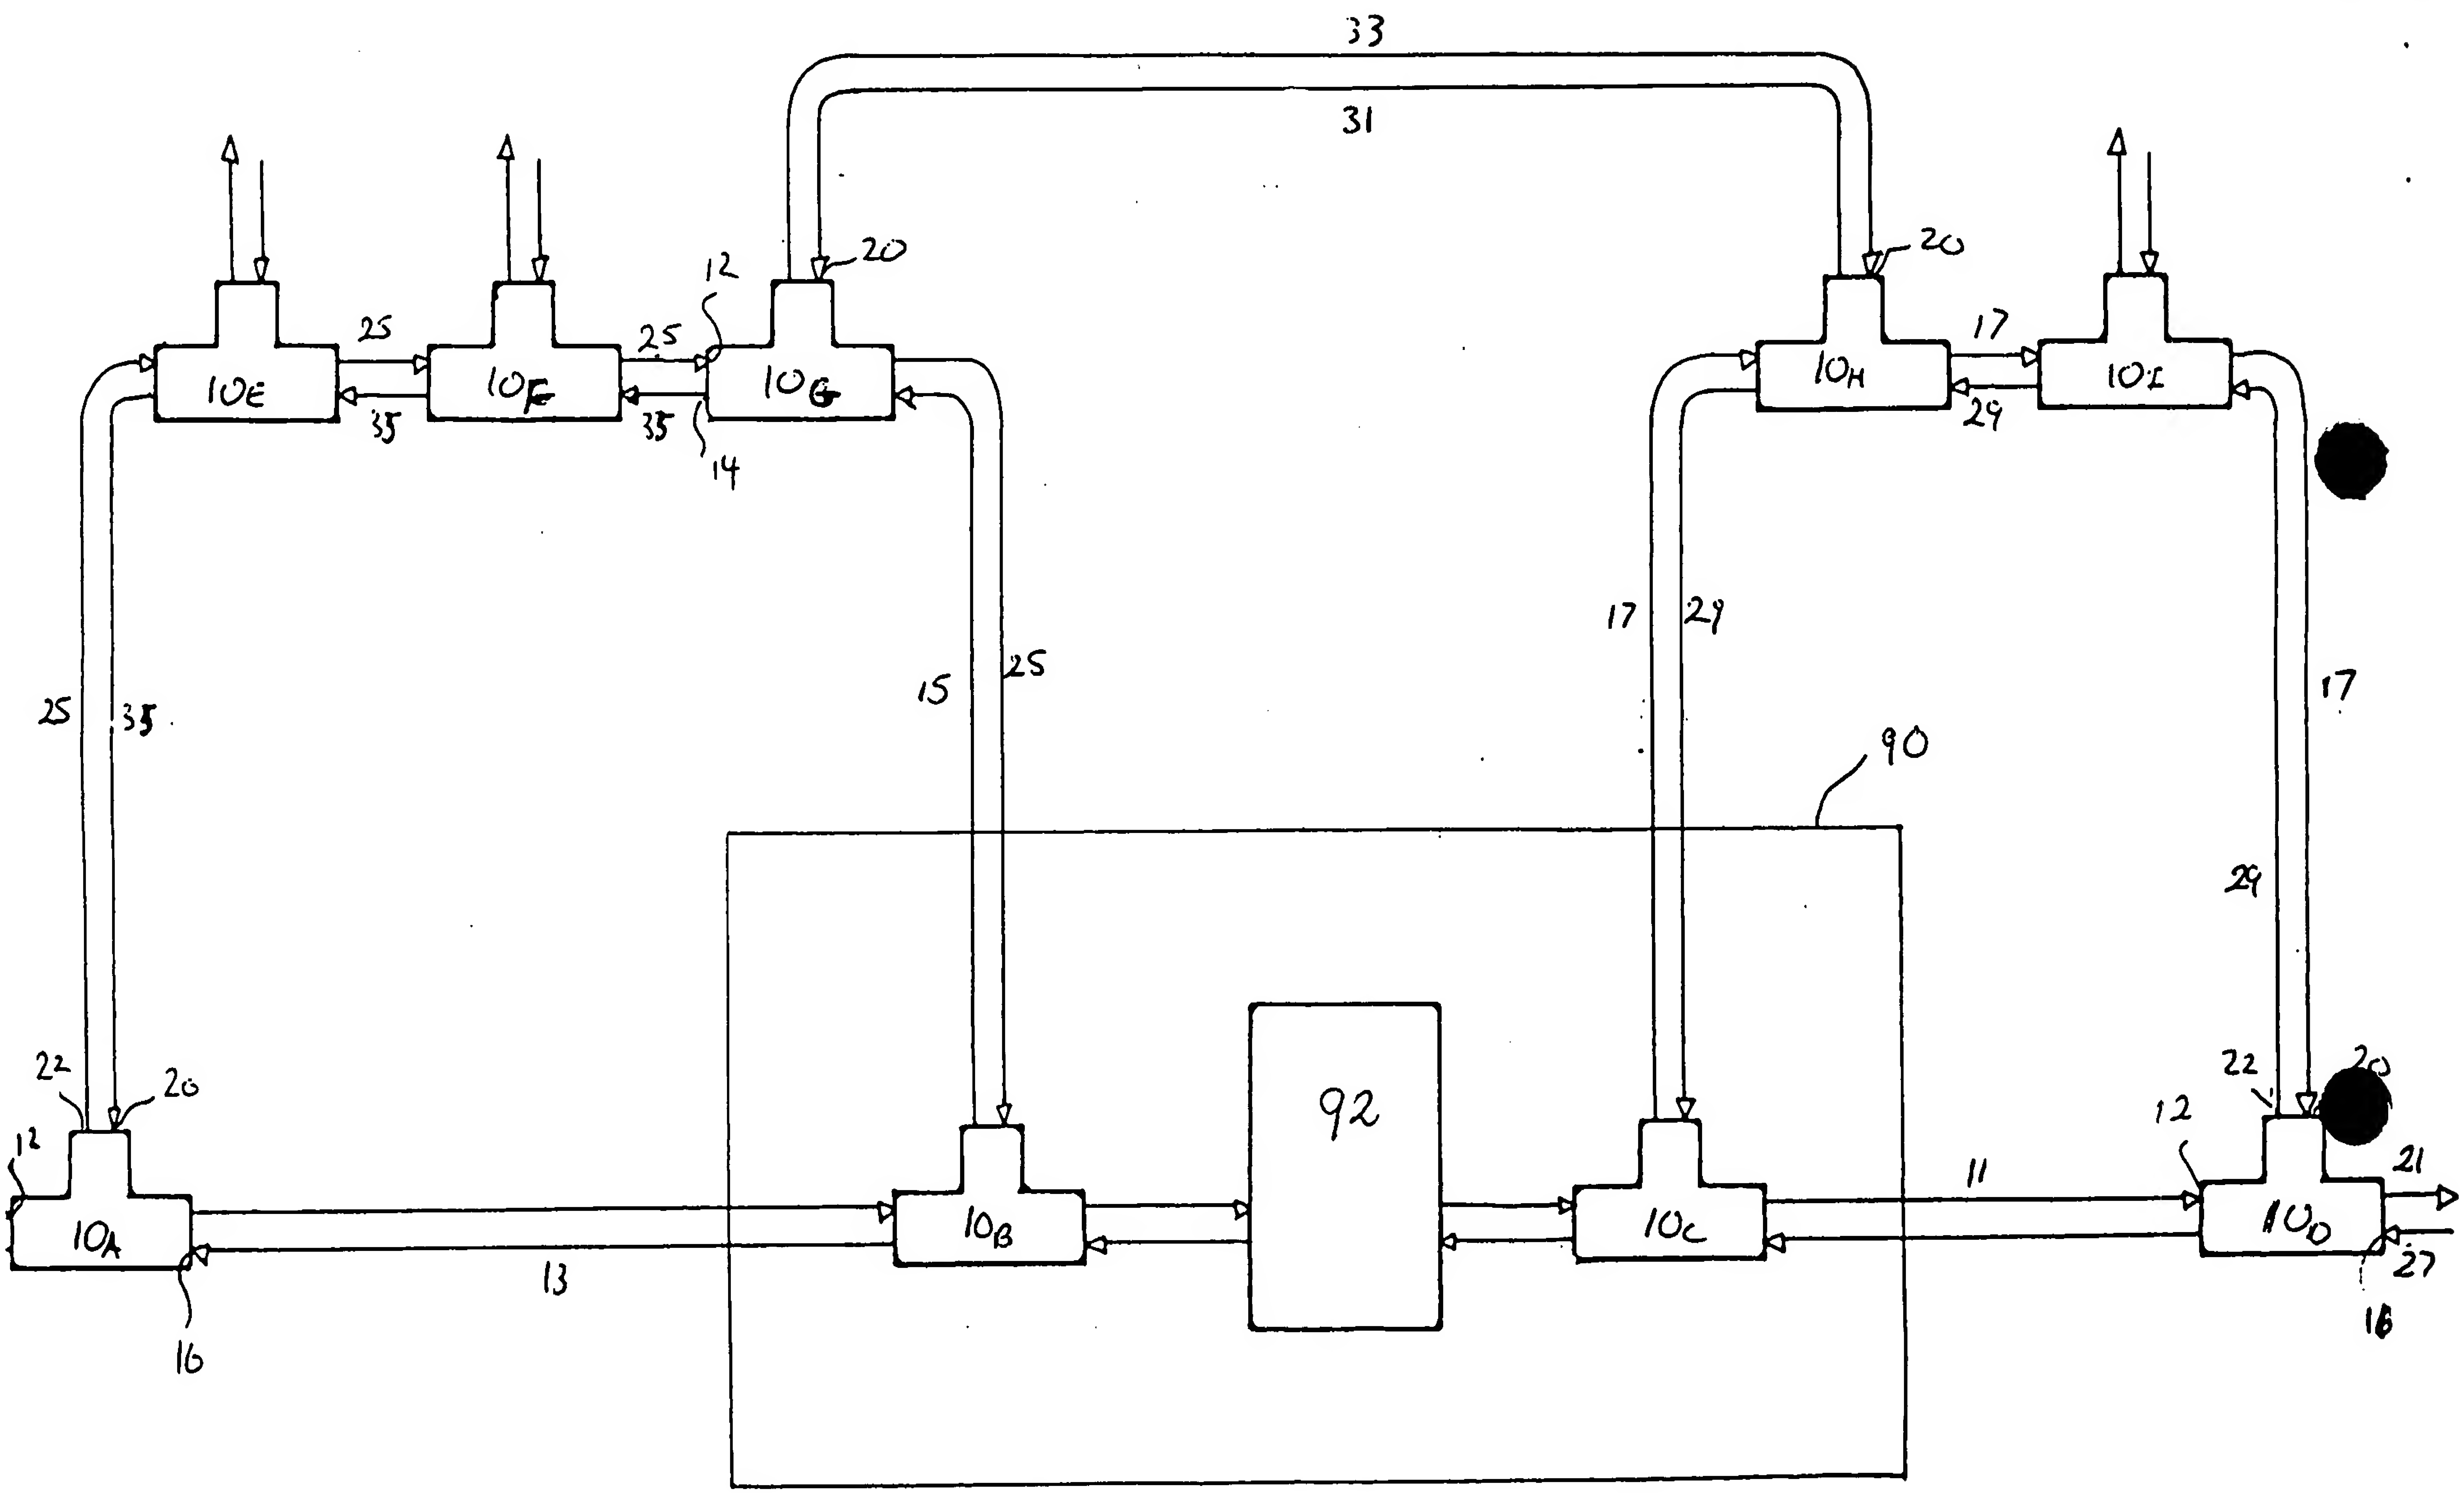

33
31
10E
10F
10G
10H
10I
10A
10B
10C
10D
92
90
25
35
12
14
20
22
16
15
17
29
11
13
21
27

FIG 16

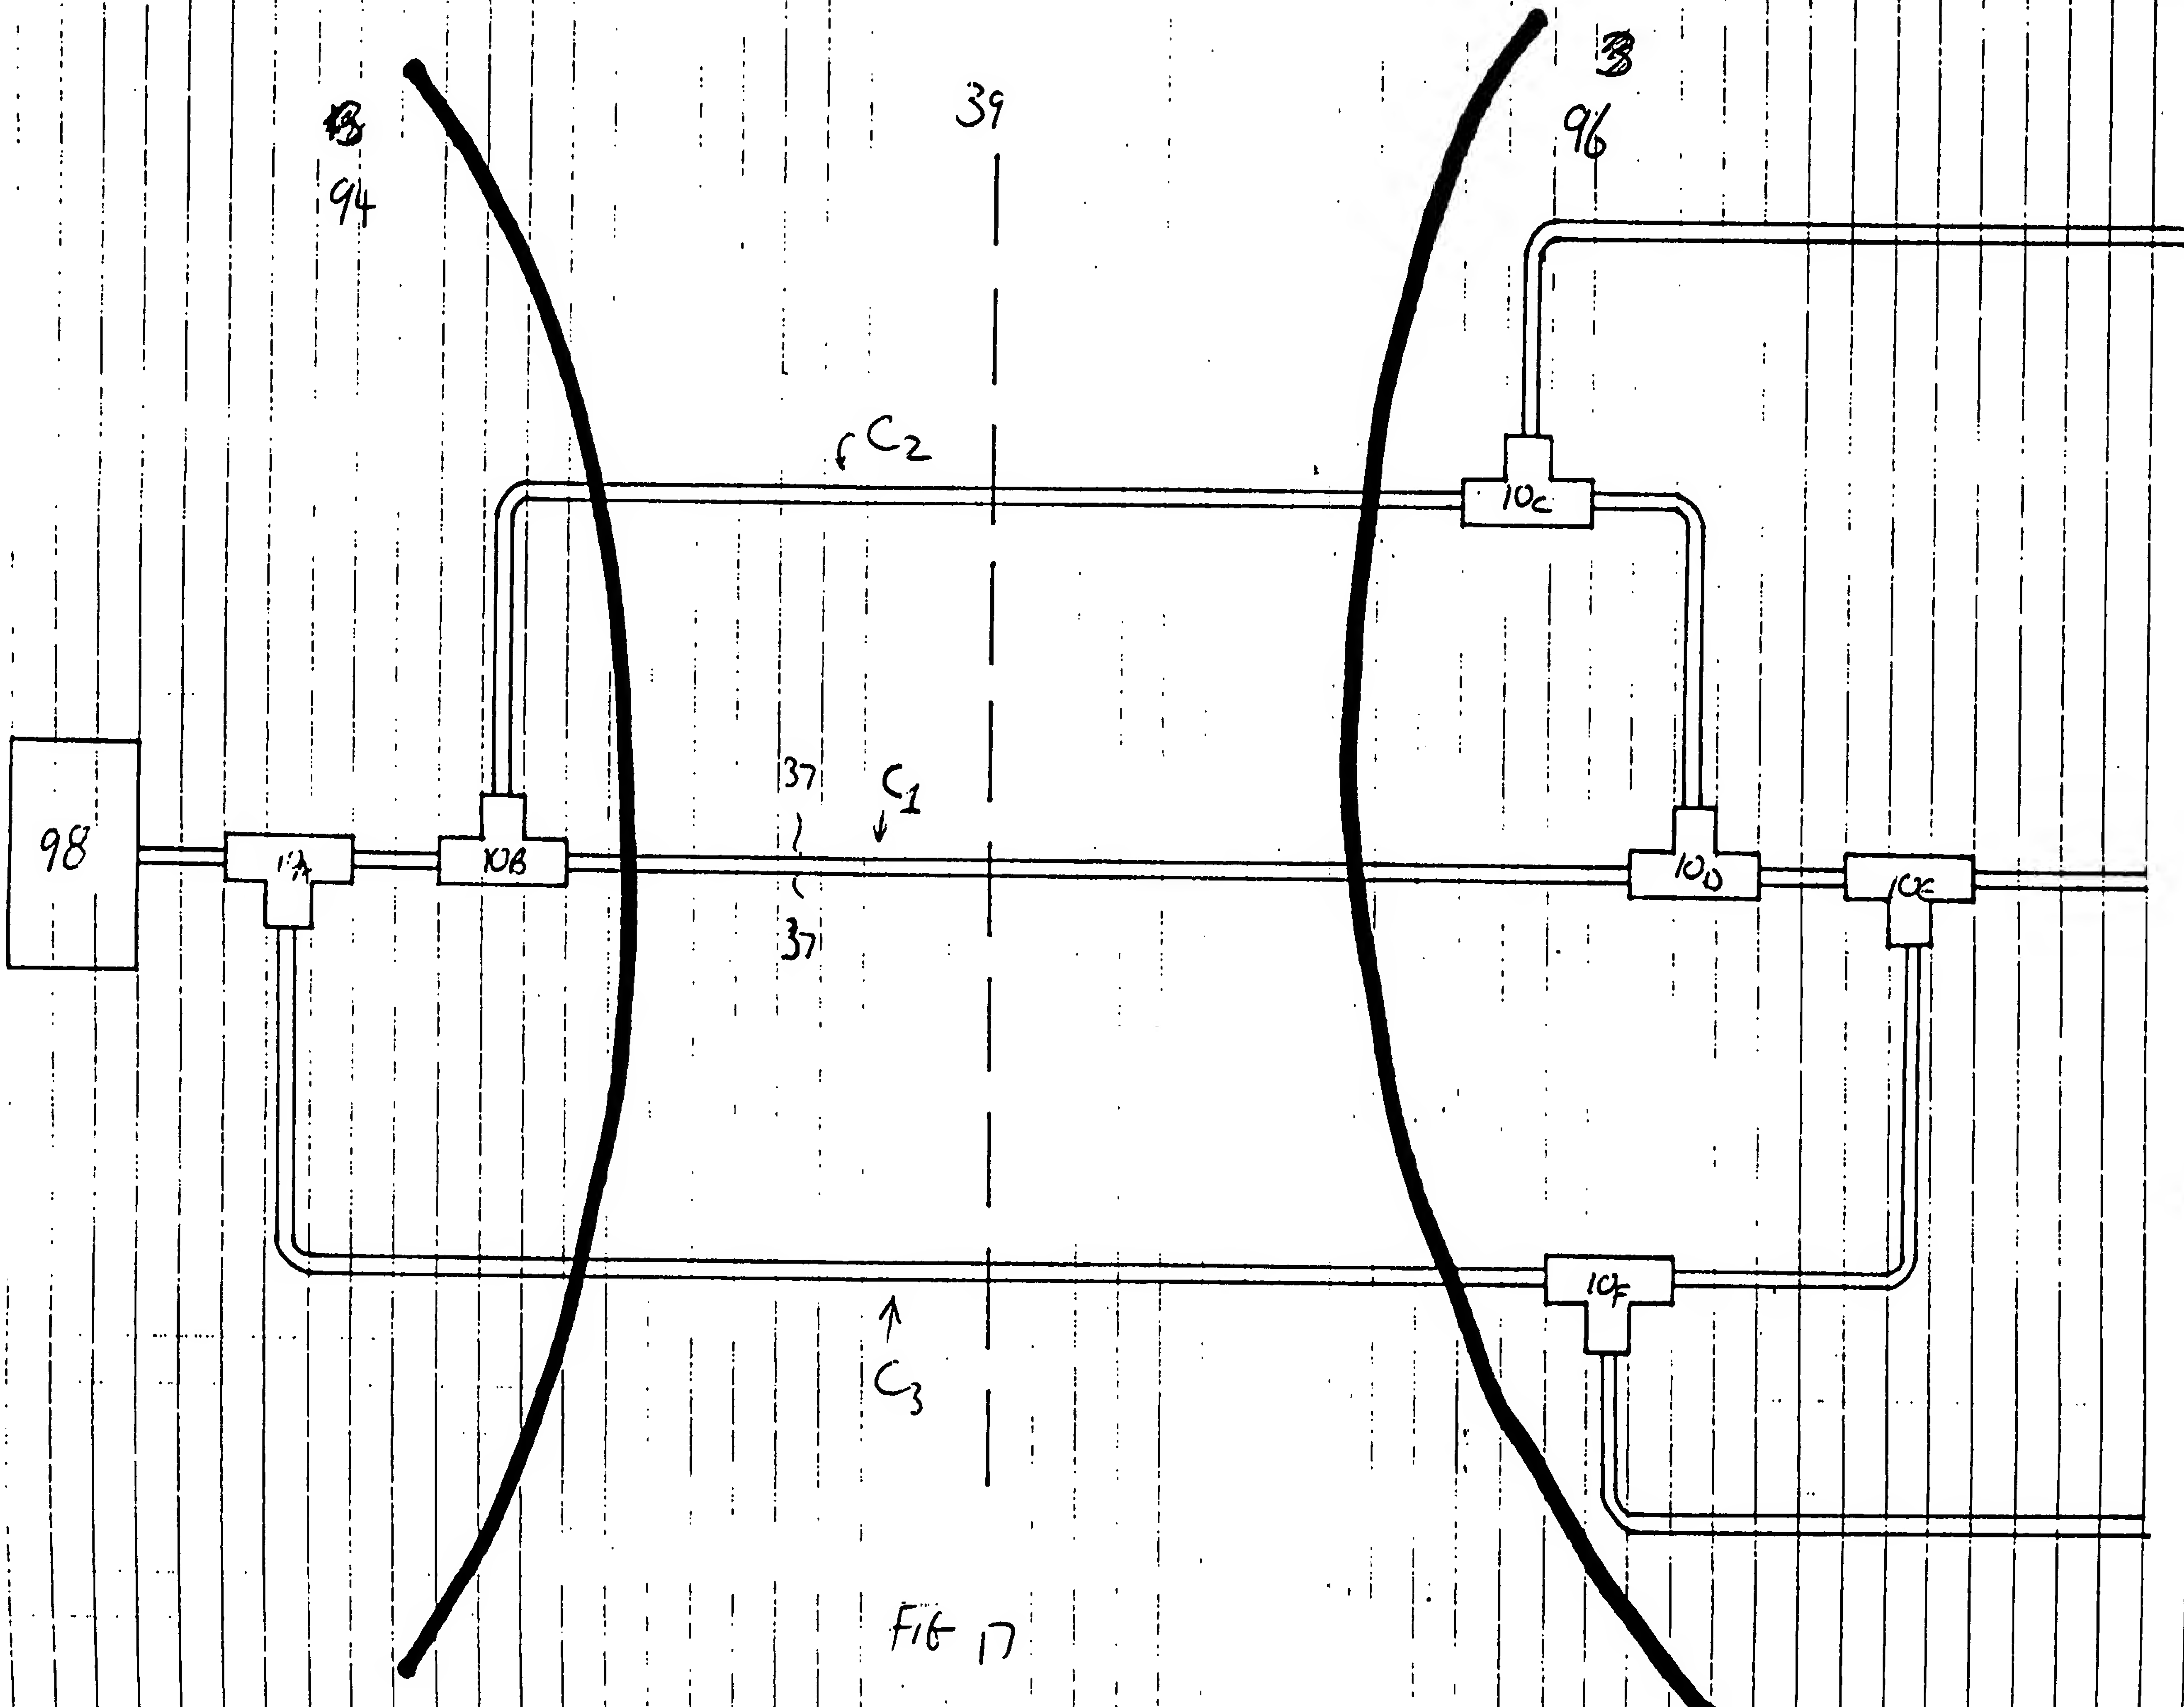

FIG 17

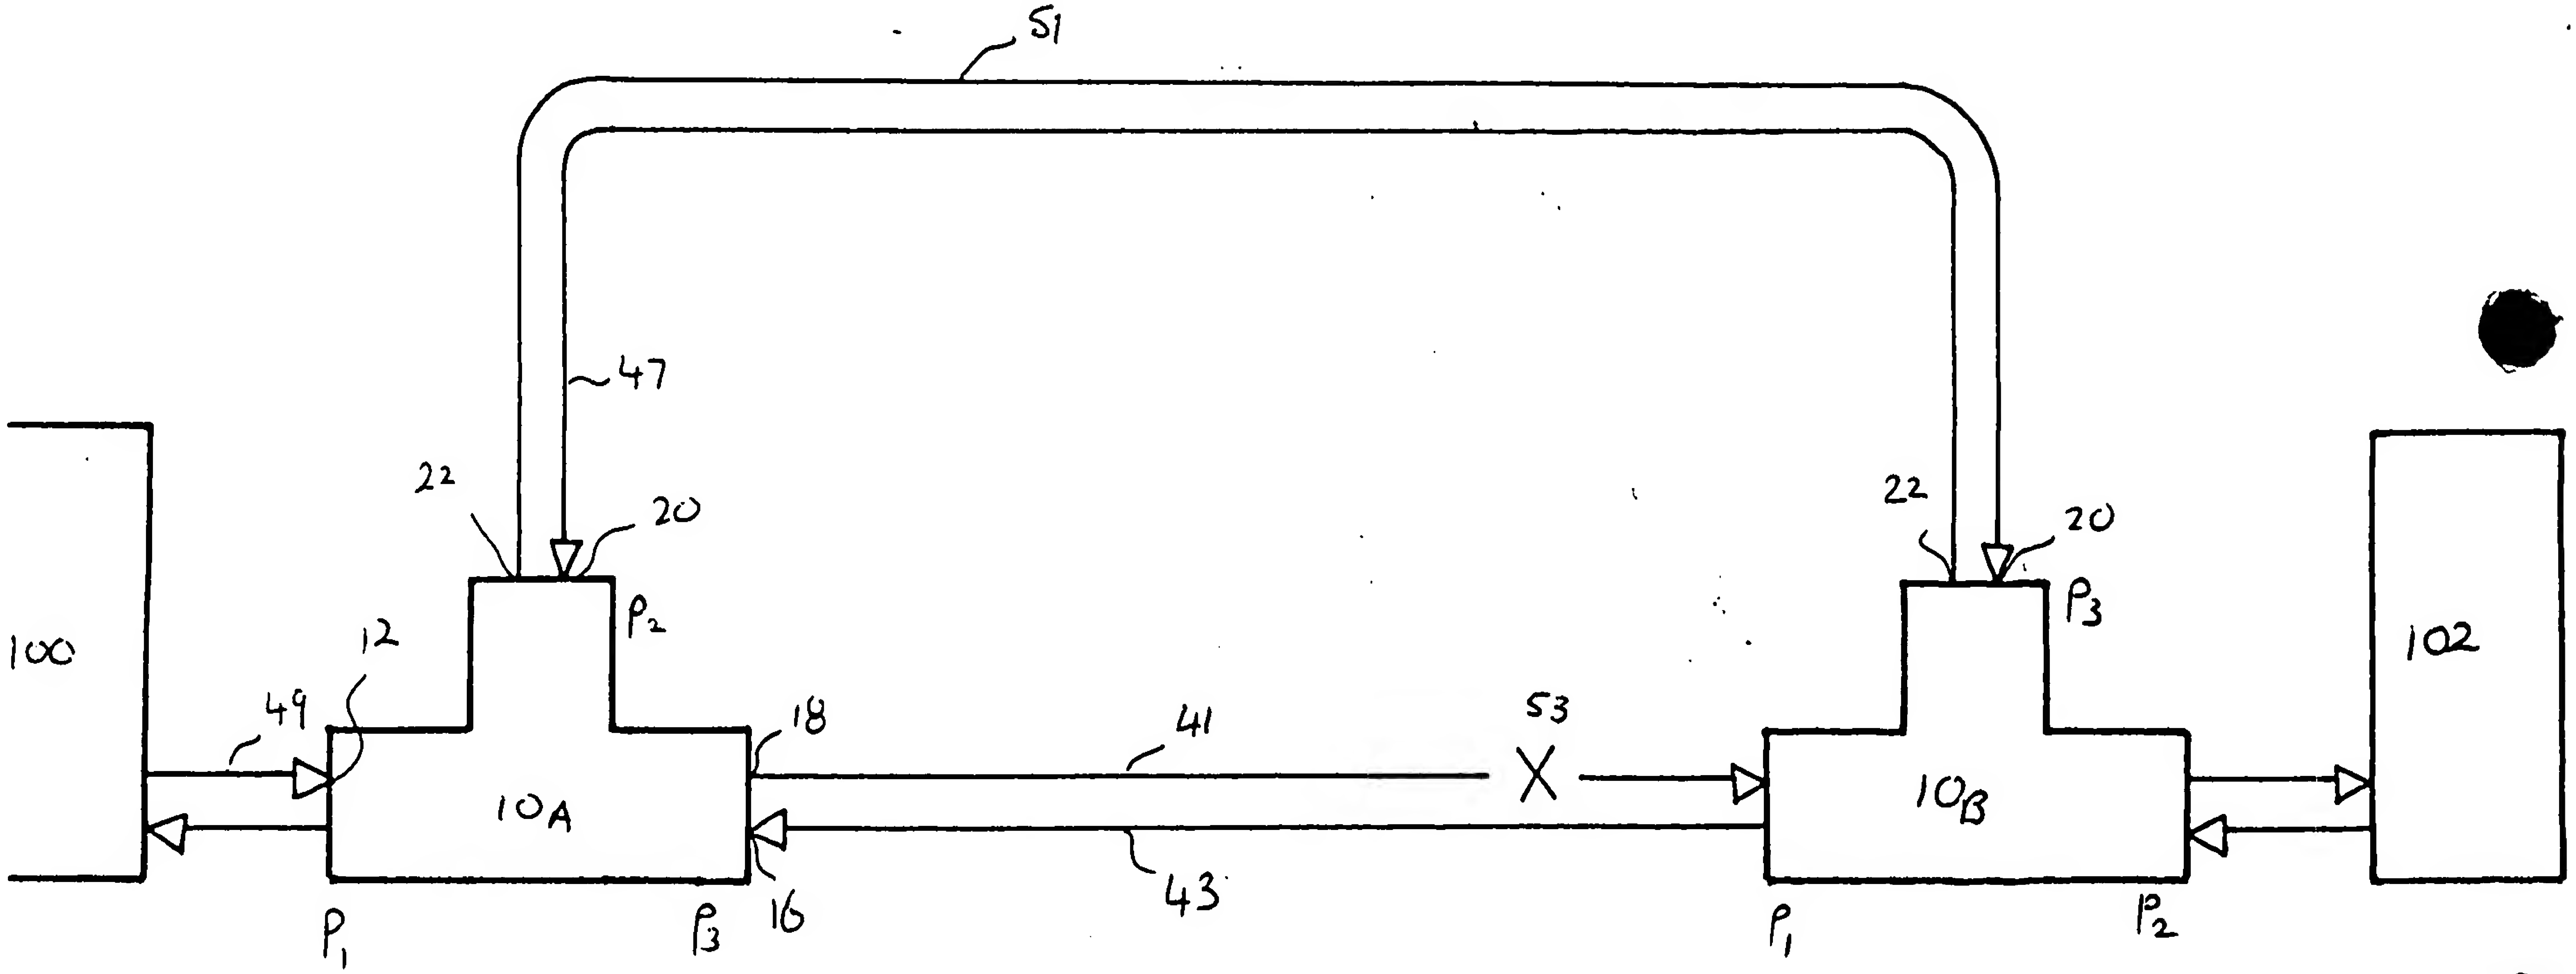

51
47
22
20
12
49
100
10A
P_1
P_2
P_3
18
16
41
43
53
22
20
P_3
10B
P_1
P_2
102

FIG18

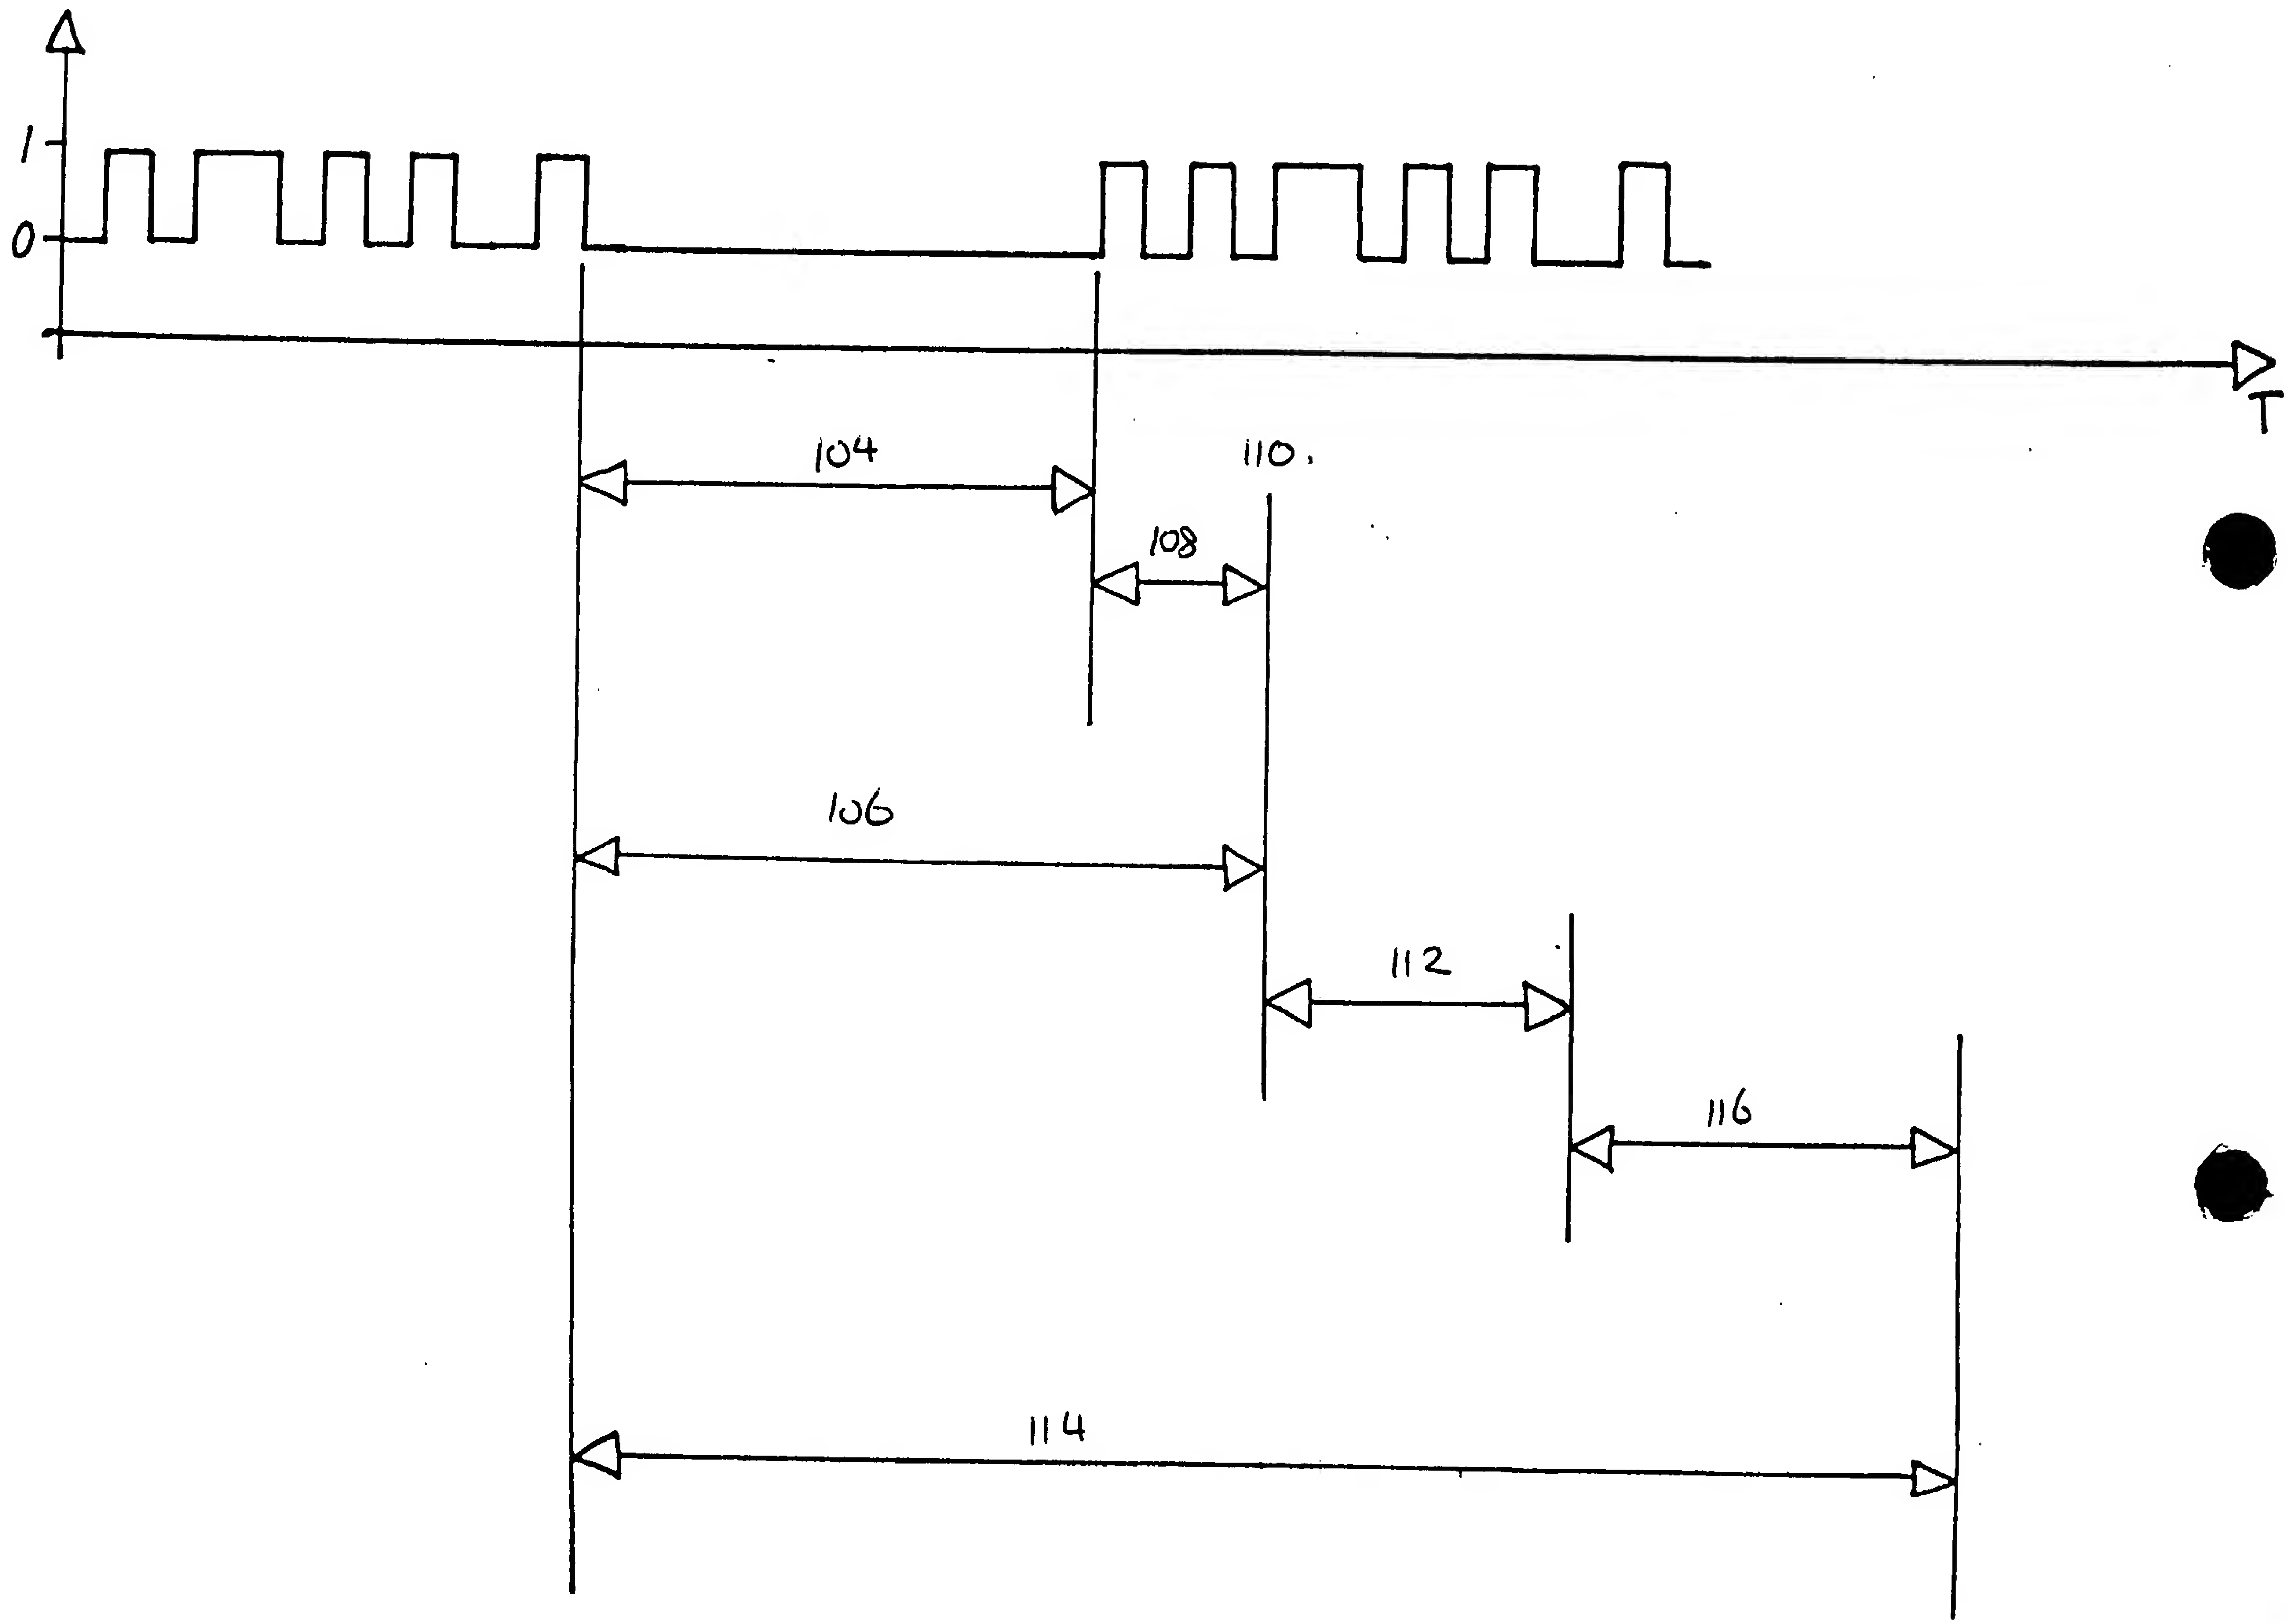

FIG 19

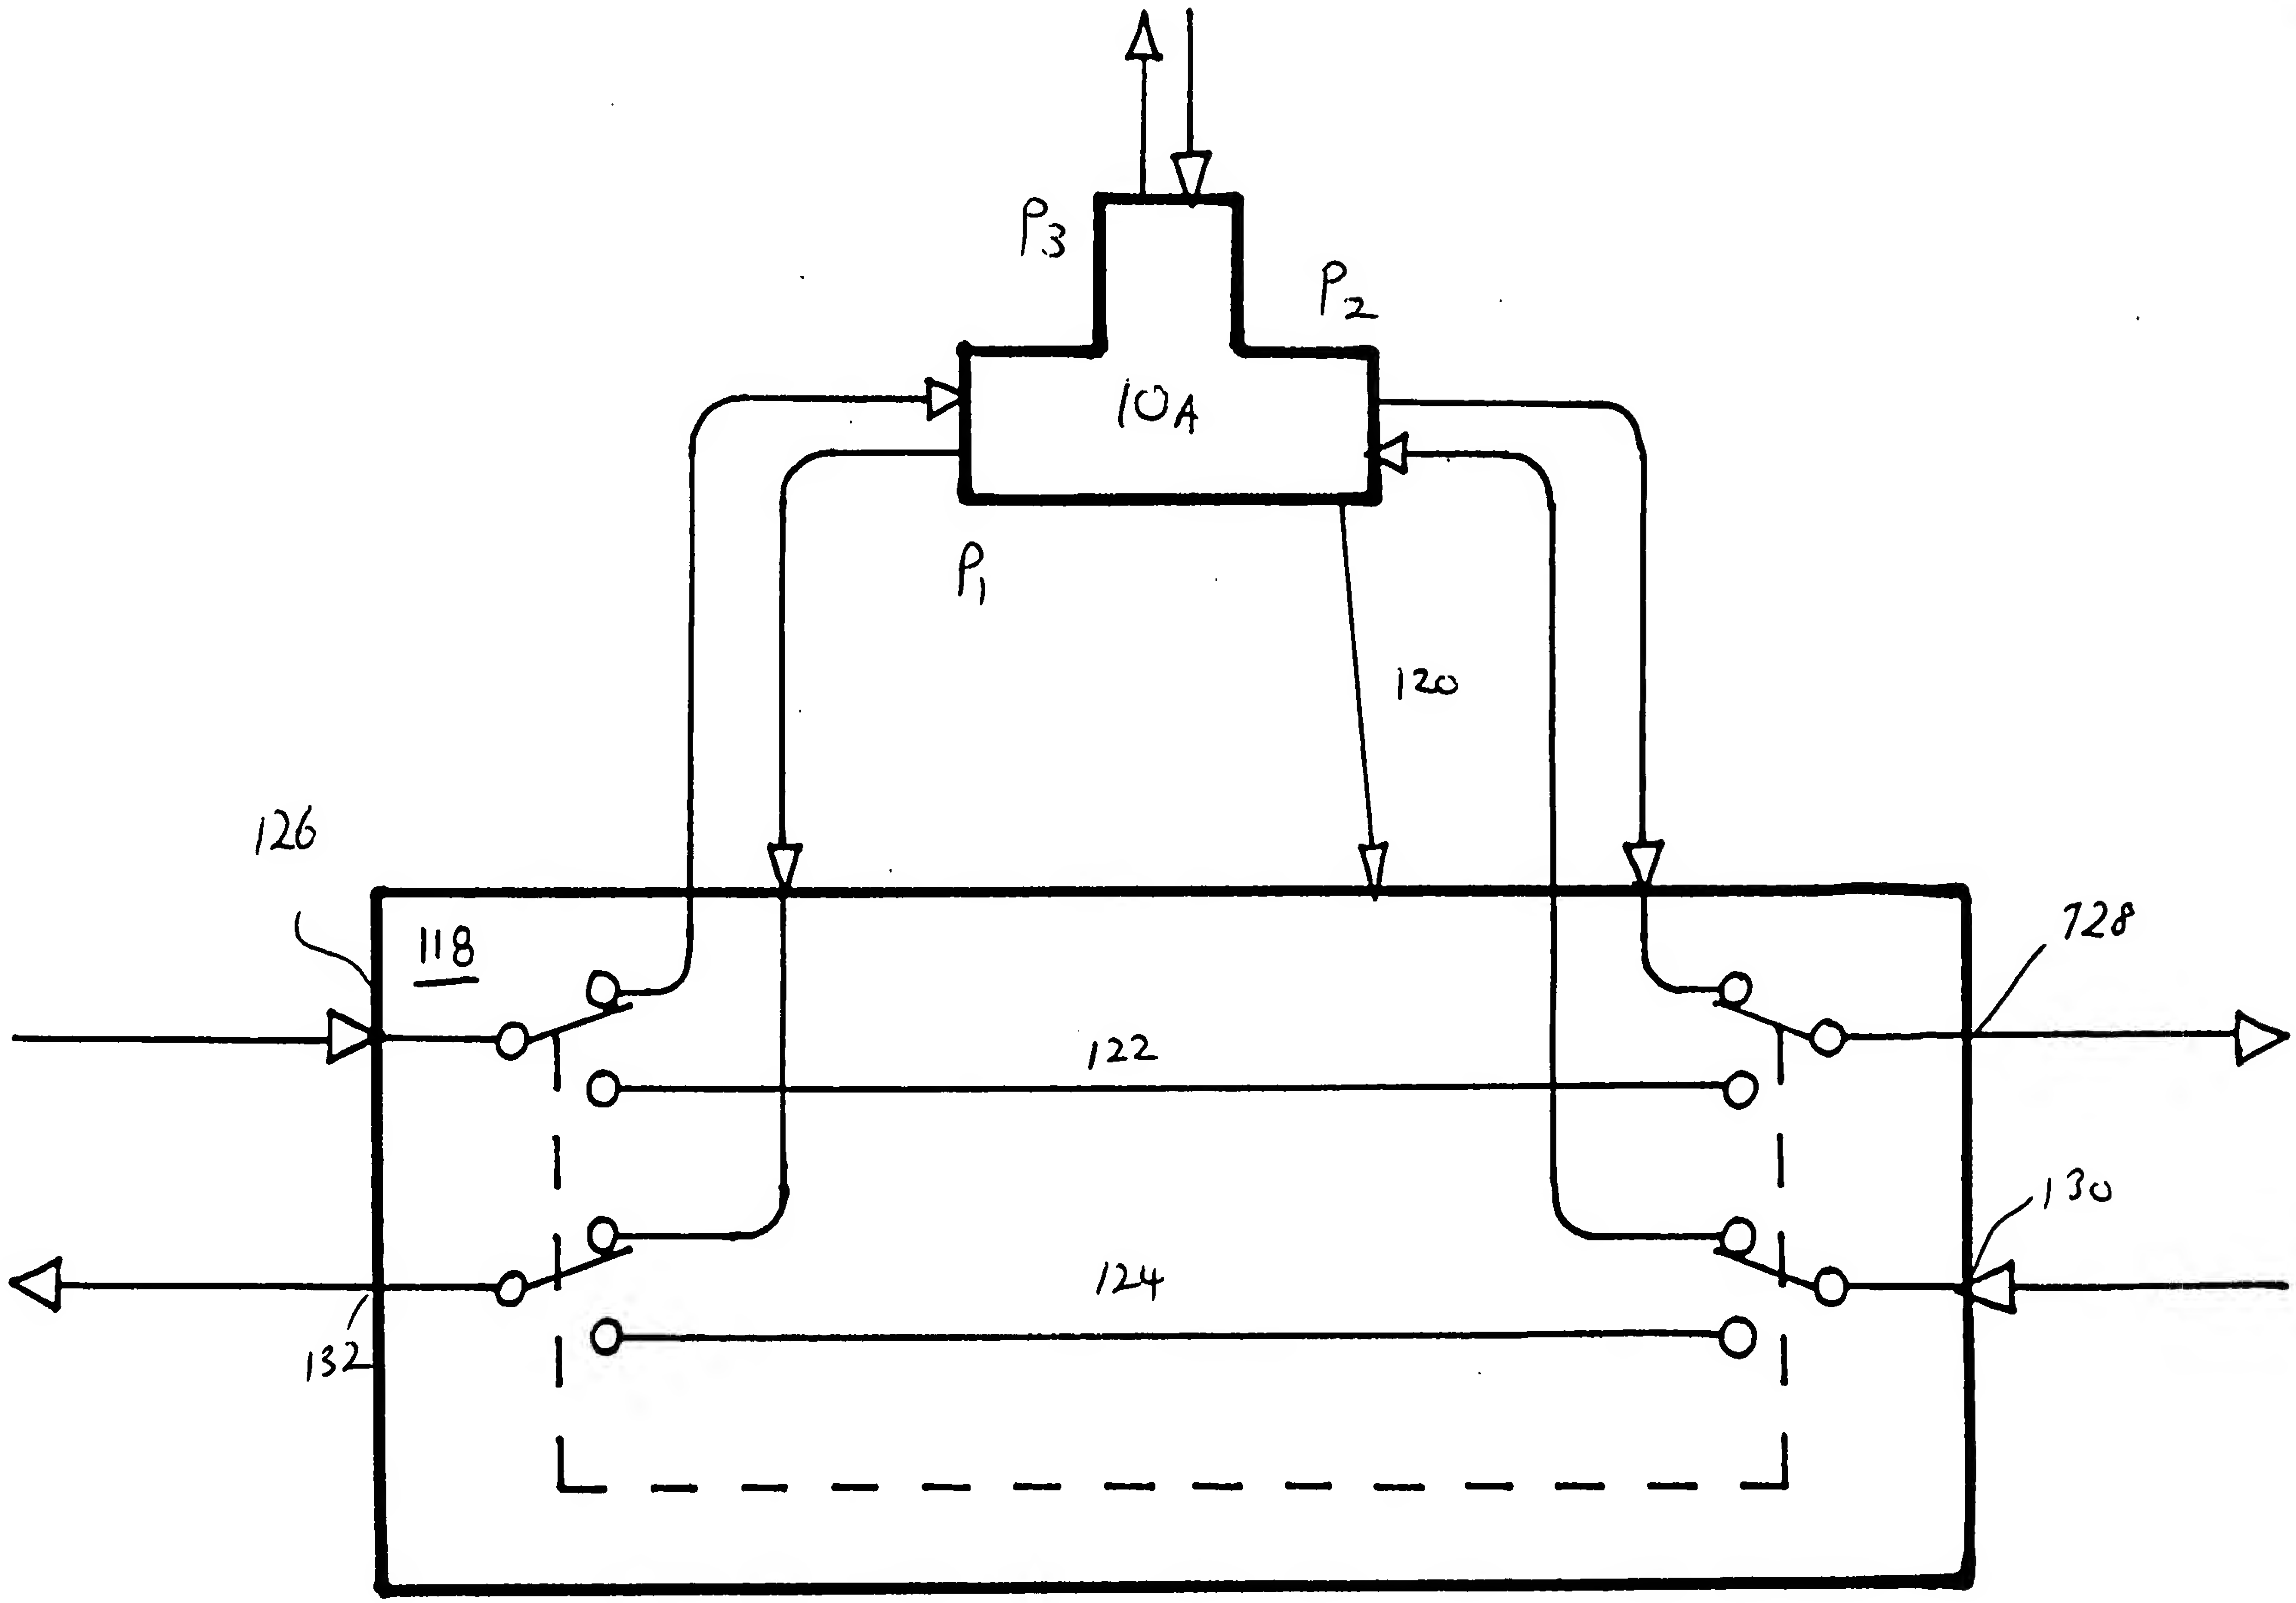

FIG 20

This Page Blank (uspto)